arXiv:1406.2587v5 [cs.SI] 5 Nov 2015

Structural Sparsity of Complex Networks: Bounded Expansion in Random Models and Real-World Graphs

Erik D. Demaine* Felix Reidl† Peter Rossmanith†
Fernando Sánchez Villaamil† Somnath Sikdar† Blair D. Sullivan‡

Abstract

This research aims to identify strong structural features of real-world complex networks, sufficient to enable a host of graph algorithms that are much more efficient than what is possible for general graphs (and currently used for network analysis). Specifically, we study the property of *bounded expansion*[1]—roughly, that any subgraph has bounded average degree after contracting disjoint bounded-diameter subgraphs. This is the strongest formalization of the well-observed notion of "sparsity" that might possibly apply to real-world networks and allow many previously intractable problems to become tractable.

On the theoretical side, we analyze many previously proposed models for random networks and characterize which ones have bounded expansion. We show that, with high probability, (1) graphs sampled with either the Molloy–Reed configuration model (including a variation of the model which achieves high clustering) or the Chung–Lu model with a prescribed sparse degree sequence (including heavy-tailed degree distributions); (2) perturbed bounded-degree graphs; (3) stochastic block models with small probabilities; result in graphs of bounded expansion. We also prove that the Kleinberg model and the Barabási–Albert model, in fairly typical setups, contain large one-subdivisions of cliques and thus do not result in graphs of bounded expansion.

On the practical side, we give experimental evidence that many complex networks have bounded expansion, by measuring the closely related "low treedepth coloring number" on a corpus of real-world data.

On the algorithmic side, we show how tools provided by the bounded expansion framework can be used to efficiently solve the following common network analysis problems: for a fixed graph H, we obtain the fastest-known algorithm for counting the number of induced H-subgraphs and the number of H-homomorphisms; and we design linear algorithms for computing several centrality measures.

Keywords: complex networks · social networks · sparsity · bounded expansion · Molloy-Reed · scale-free · treedepth · parameterized algorithms

*Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA. edemaine@mit.edu.

†Theoretical Computer Science, Dept. of Computer Science, RWTH Aachen University, Aachen, Germany, {reidl,rossmani,fernando.sanchez,sikdar}@cs.rwth-aachen.de.

‡Department of Computer Science, North Carolina State University, Raleigh, NC, blair_sullivan@ncsu.edu.

[1]Not to be confused with the notion of expansion related to expander graphs.

1 Introduction

Complex networks vs. structural graph algorithms. Social networks (such as Facebook or physical disease propagation networks), biological networks (such as gene interactions or brain networks), and informatics networks (such as autonomous systems) are all examples of *complex networks*, which have been the attention of much study in recent years, given the surge of available network data. Fed as graphs, complex networks seem to share several structural properties. Perhaps most famous is the *small-world* property ("six degrees of separation"): typical distances between vertex pairs are small compared to the size of the network. Another important property is that their degree distribution tends to be heavy-tailed, i.e. distributions whose tails are not exponentially bounded. In many cases, this degree distribution is close to a *power law*: the fraction of vertices of degree k is proportional to $k^{-\gamma}$, for some constant γ typically between 2 and 3. Networks furthermore often exhibit *high clustering* and admit a natural division into a *community structure*. Despite the clustering property, complex networks are *sparse*: the ratio of edges to vertices (edge density) is usually small.

On the other hand, the field of structural graph algorithms has led to impressively efficient and precise algorithms (efficient PTASs, subexponential fixed-parameter algorithms, linear kernelizations, etc.) for increasingly general families of graphs; see, e.g., [25–28, 32, 36, 47]. Many such results proved initially for planar graphs have since been extended to bounded-genus graphs, graphs of bounded local treewidth, and graphs excluding a fixed minor; yet such results are known to be impossible for general graphs. Can we apply these powerful algorithms to analyze complex networks?

We propose the following litmus test for whether a type of network sparsity is "useful": Does it enable efficient algorithms for a broad set of NP-hard problems? Unfortunately, the above-mentioned structural properties of complex networks seem too weak to enable better algorithms, while the discussed graph classes seem too restrictive to apply to complex networks. The goal of this paper is to bridge this gap, by identifying a more general graph class that simultaneously enables better algorithms and includes many complex networks.

Bounded expansion. In general, complex networks seem to exhibit an intermediate-scale structure composed of small dense parts—representing clusters/communities—that are sparsely interconnected. This hierarchical behavior has been established for many networks [58] and is consistent with the tree-like intermediate structure observed in [4].

How can this notion be captured formally? If we contract disjoint small-diameter subgraphs (representing potentially dense local substructures in the

network), then the resulting graph (representing the global connectivity of these substructures) should be sparse. This is the notion of an *r-shallow minor*, where r is the maximum diameter of the subgraphs that were contracted in the construction process. (For formal definitions, see Section 2.) We cannot expect the edge density of all r-shallow minors to be constant (then r would play no role), but we require it to grow as any function of r, and thus be independent of the size of the graph. A graph class for which this property holds has *bounded expansion*, a concept introduced by Nešetřil and Ossona de Mendez [71].

Theoretical results. Since the definition of bounded expansion applies to graph classes instead of individual graphs, it is impossible to settle this question empirically. To ground our hypothesis in theory, we analyze several random graph models which were designed to mimic the properties of specific types of real world networks. Although not perfect, random graph models play a central role in network analysis, both to guide our understanding of complex networks and as a convenient source of synthetic data for algorithm testing and validation. In our case, random graph models allow us to determine whether (synthetic) complex networks have bounded expansion with high probability. We analyze several popular random graph models:

(i) the configuration model [69] and the Chung–Lu model [21, 22] with specified asymptotic degree sequences, which includes graphs with heavy-tailed degree distributions;

(ii) a variant of the configuration model which achieves high clustering [9];

(iii) a significant generalization of Erdős–Rényi graphs we call *perturbed bounded-degree graphs* (allowing the network to be built on top of a fixed or random base graph of bounded degree), which includes the stochastic block model with small probabilities;

(iv) the Kleinberg model [48, 49];

(v) and the Barabási–Albert model [6, 10].

We will show that the configuration model, the Chung–Lu model and perturbed bounded-degree graphs have bounded expansion, while the Kleinberg model and the Barabási–Albert model do not (actually they are not even *nowhere dense*, a strict generalization of bounded expansion).

Experimental results. We present an experimental study suggesting that important real-world networks have small *grad*, which is the density measure for single graphs that defines the expansion for graph classes. Interestingly, the algorithmic tools that become efficient when the grad of a graph is small can be directly applied without knowing its actual value.

We will make extensive use of *low-treedepth colorings*: for any integer p, a graph can be colored with $f(p)$ colors such that any set of at most $p-1$ colors induced a graph of low treedepth, where f only depends on the grad of the graph. Generally the running time of algorithms based on low-treedepth colorings depends heavily on the number of colors $f(p)$. In Section 5 we present experimental results obtained by computing and evaluating low-treedepth colorings with a simple algorithm in a variety of real-world networks.

Our results show that, in general, real networks exhibit even better structure (require fewer colors for a p-centered coloring) than randomly generated graphs with the same degree distribution via the configuration model. These results support our hypothesis that "community structure" (not captured by the degree distribution) further increases the algorithmic tractability.

Algorithmic results. With both theoretical and empirical results in hand, we exploit the aforementioned tools for graphs with small grad in Section 6 to solve typical problems for complex networks: First we develop a faster algorithm than the one presented in [73] to count subgraph homomorphisms based on low treedepth colorings. Counting subgraphs is fundamental to *motif counting*, a widely used concept to analyze networks. Then we develop an algorithm which computes localized versions of several centrality measures in linear time on graphs of bounded expansion. Specifically, we present:

i. A linear-time algorithm to count the number of times a fixed subgraph appears as an (induced) subgraph/homomorphism in graphs of bounded grad. We do this by improving the previously best known algorithm to count the appearances of a structure of size h on graphs of treedepth t from $O(2^{ht}ht\cdot n)$ to $O(t^h 6^h h^2\cdot n)$, thus removing the exponential dependency on t, while keeping the algorithm simple and avoiding big hidden constants.

ii. A linear-time algorithm to compute localized variants (i.e., computed in a constant-radius neighborhood around each vertex) of the *closeness centrality* and two other related measures. The constant in the running time depends on the radius and the grad of the graph.

For the second algorithm we provide experimental results which indicate that the local variants of these centrality measures can be used to estimate the top 10 percent of the most central nodes.

Previous results. There is substantial empirical work studying structural properties of complex networks, so we focus here on work closest to structural graph theory. In general, large real-world complex networks are not easily classified as either low- or high-dimensional. In particular, data-mining

tools which implicitly assume low dimensionality (such as singular value decomposition) produce models and results incompatible with observed structure and dynamics, yet traditional high-dimensional tools (like sampling) often fail to achieve measure concentration due to the extreme sparsity of the networks. Adcock et al. [4] recently empirically established that, when compared with a suite of idealized graphs[2], realistic large social and informatics networks exhibit meaningful "tree-like" structure at an intermediate scale. Their work related the k-core structure (whose extremal statistic is the degeneracy) to the networks' Gromov hyperbolicity and tree decompositions. Unfortunately, they showed that straightforward applications of these measures are often insufficient to reveal meaningful structure because of noisy/random edges in the network which contradict the strict structural requirements. (For example, several families of popular random graph models have been shown to have very large treewidth [37].)

Some simple preliminary observations about networks and bounded expansion were made in [35], such as the *linear growth model* not having bounded expansion since it was known that it contains growing bi-cliques, and they conjecture that Barabási-Albert is somewhere dense a.a.s. Here we prove that it is at least somewhere dense with non-vanishing probability. For the *random intersection graph* model, which is used to model real world networks where connections depend on shared attributes, it was shown in [31] that it has bounded expansion exactly when it is degenerate.

2 Preliminaries

For a natural number n we use the notation $[n] := \{1, \ldots, n\}$. For a graph G, we denote by $\Delta(G)$ its maximal degree and by $\omega(G)$ its clique number. We will make use of the following graph operations. For graphs G_1, G_2, the *complete join* $G_1 * G_2$ is the graph obtained by first taking the disjoint union of G_1, G_2 and then connecting every vertex of $V(G_1)$ to every vertex $V(G_2)$. For example, $G * \overline{K}_2$ is the graph obtained from G by adding two universal vertices. The *lexicographic product* $G_1 \bullet G_2$ is the graph with vertices $V(G_1) \times V(G_2)$ and edges

$$(u, x)(v, y) \in E(G_1 \bullet G_2) \iff uv \in E(G_1) \text{ or } (u = v \text{ and } xy \in E(G_2)).$$

[2]representing low-dimensional structures, common hierarchical models, constant-degree expanders, etc.

2.1 Bounded expansion classes

We will use $(\leqslant r)$-*subdivisions* to formalize the notion of shallow topological minors. A $(\leqslant r)$-subdivision of a graph M is any graph which can be obtained from M by replacing edges with disjoint paths of length at most $r+1$.

Definition 1 (Shallow topological minor [71])**.** *A graph M is an* r-shallow topological minor *if a $(\leqslant 2r)$-subdivision of M is isomorphic to a subgraph G' of G. We call G' a* model of M in G*. For simplicity, we assume by default that $V(M) \subseteq V(G')$ such that the isomorphism between M and G' is the identity when restricted to $V(M)$. The vertices $V(M)$ are called* nails *and the vertices $V(G') \setminus V(M)$* subdivision vertices*. The set of all r-shallow topological minors of a graph G is denoted by $G \mathbin{\widetilde{\nabla}} r$.*

The theory of bounded expansion classes crucially relies on a 'parameterization' of shallow minors. The *grad* of a graph is one such parameterization:

Definition 2 (Topological grad)**.** *For a graph G and an integer $r \geqslant 0$, the* topological greatest reduced average density (topological grad) at depth r *is defined as*

$$\widetilde{\nabla}_r(G) = \max_{H \in G \mathbin{\widetilde{\nabla}} r} \frac{|E(H)|}{|V(H)|}.$$

For a graph class $\mathcal{G}$, define $\widetilde{\nabla}_r(\mathcal{G}) = \sup_{G \in \mathcal{G}} \widetilde{\nabla}_r(G)$.

Given the notion of shallow topological minors and topological grad, we can now define what it means for a class to have bounded expansion.

Definition 3 (Bounded expansion)**.** *A graph class $\mathcal{G}$ has* bounded expansion *if and only if there exists a function f such that for all $r \geqslant 0$, it holds that $\widetilde{\nabla}_r(\mathcal{G}) < f(r)$.*

When introduced in [71], bounded expansion was originally defined using a characterization based on the notion of *shallow minors*: H is an r-shallow minor of G if H can be obtained from G by contracting disjoint r-balls and then taking a subgraph. Taking the maximum over the density of all r-shallow minors then defines the *grad*. An r-ball in a graph G is a subgraph $G' \subseteq G$ with the property that there exists $v \in V(G')$ such that for all $u \in V(G')$, $d_{G'}(u, v) \leqslant r$. Both characterizations are equivalent, hence we will for brevity's sake simply talk about the grad.

We note that graphs excluding a topological minor—in particular planar graphs and bounded-degree graphs—have bounded expansion. This generalizes to graphs excluding a minor (and thus to those of bounded treewidth). Finally, we point out that bounded expansion implies bounded degeneracy,

where a graph G is d-degenerate if any subgraph of G contains a node of degree smaller than d. The converse does not hold.

The following alternative characterization of bounded expansion uses a special coloring number with nice algorithmic properties.

Definition 4 (p-centered coloring [71])**.** *Given a graph G, let $c\colon V(G) \to [r]$ be a vertex coloring of G with r colors. We say that the coloring c is p-*centered*, for $p \geqslant 2$, if any connected subgraph of G either receives at least p colors or contains some color exactly once. Define $\chi_p(G)$ to be the minimum number of colors needed for a $(p+1)$-centered coloring.*

While this definition looks rather cryptic, it is easy to see that every graph has a p-centered coloring for any p: simply assign a distinct color to each vertex of the graph. Note that p-centered colorings are proper colorings for $p \geqslant 2$ and in particular, χ_1 is precisely the chromatic number. Typically, the number of colors q is much larger than p and one is interested in minimizing q.

The following structural property, which follows directly from the equivalence between centered colorings and *treedepth*, make them an attractive tool for algorithm design.

Proposition 1 (Low treedepth colorings [72])**.** *Let $\mathcal{G}$ be a graph class of bounded expansion. There exists a function f such that for every $G \in \mathcal{G}$, $r \in \mathbf{N}$, the graph G can be colored with $f(r)$ colors so that any $i < r$ color classes induce a graph of treedepth $\leqslant i$ in G. This coloring can be computed in linear time.*

For those unfamiliar with this notion, the *treedepth* of a graph G is the lowest rooted forest whose closure contains G as a subgraph. A *treedepth decomposition* of G is simply a forest with vertex set $V(G)$ witnessing this fact. To put this width measure into perspective: a graph of treedepth at most t cannot contain a path of length 2^t *and* has pathwidth at most $t-1$. An example of a 5-centered coloring and treedepth-decompositions of a selected subset of colors can be found in Figure 2 on page 42. For more information about treedepth see e.g. [87].

Nešetřil and Ossona de Mendez show that graph classes of bounded expansion are precisely those for which there exists a function f such that every member G of the graph class satisfies $\chi_p(G) \leqslant f(p)$ (see Theorem 7.1 in [71]). In [72], the authors also showed how to obtain a p-centered coloring with at most $P(f(p))$ colors for each fixed p in linear time, where P is some polynomial of degree roughly 2^{2^p}. We will make use of this algorithm in Sections 5 and 6 and see that the actual number of colors is manageable.

When working with random graphs, we will make heavy use of the following alternative characterization of graphs of bounded expansion:

Proposition 2 (Nešetřil, Ossona de Mendez, Wood [73, 74])**.** *A class $\mathcal{C}$ of graphs has bounded expansion if and only if there exist real-valued functions $f_{thresh}, f_{\deg}, f_{\widetilde{\nabla}}, f_H \colon \mathbf{R}^+ \to \mathbf{R}$ such that the following two conditions hold:*

i. For all $\varepsilon > 0$ and for all $G \in \mathcal{C}$ with $|G| > f_{thresh}(\varepsilon)$, it holds that

$$\frac{1}{|V(G)|} \cdot |\{v \in V(G) \colon \deg(v) \geqslant f_{\deg}(\varepsilon)\}| \leqslant \varepsilon.$$

ii. For all $r \in \mathbf{N}$ and for all $H \subseteq G \in \mathcal{C}$ with $\widetilde{\nabla}_r(H) > f_{\widetilde{\nabla}}(r)$, it follows that

$$|V(H)| \geqslant f_H(r) \cdot |V(G)|.$$

Intuitively, Proposition 2 characterizes classes of graphs with bounded expansion as those where:

i. all sufficiently large members of the class have a small fraction of vertices of large degree;

ii. all subgraphs of $G \in \mathcal{C}$ whose shallow topological minors are sufficiently dense must necessarily span a large fraction of the vertices of G.

Finding the first pair of functions f_{thresh},$f_{\deg}$ is usually straightforward, the real challenge lies in proving that functions $f_{\widetilde{\nabla}}$,f_H exist. Consider the contrapositive of the second condition: in terms of random graph models, we want to show that subgraphs that span at most a $f_H(r)$-fraction of the vertices of G have their (topological) grad bounded by $f_{\widetilde{\nabla}}$ with high probability.

One of the useful properties of bounded expansion classes is that their expansion does not increase arbitrarily through the addition of a universal vertex.

Lemma 1. *For every graph G it holds that*

$$\frac{\widetilde{\nabla}_r(G)}{\widetilde{\nabla}_r(G)+1}(\widetilde{\nabla}_r(G)+1) < \widetilde{\nabla}_r(G * K_1) < \widetilde{\nabla}_r(G)+1$$

Proof. It is easy to see that if H is an r-shallow topological minor of G, then $H * K_1$ is a r-shallow topological minor of $G * K_1$. The density of $H * K_1$ is then given by

$$\frac{|E(H * K_1)|}{|V(H * K_1)|} = \frac{|E(H)| + |V(H)|}{|V(H)|+1}$$

And therefore we obtain upper and lower bounds via

$$\frac{1}{1+(1/|V(H)|)}\left(\frac{|E(H)|}{|V(H)|}+1\right) \leqslant \frac{|E(H * K_1)|}{|V(H * K_1)|} < \frac{|E(H)|}{|V(H)|}+1$$

Observing that $|V(H)| \geqslant \widetilde{\nabla}_r(G)$ proves the claim. □

Therefore even a constant number of universal vertices will not influence the grad too much. In terms of graph classes this means that if $\mathcal{G}$ has bounded expansion, then the class $\{G * K_c\}_{G\in\mathcal{G}}$ for any constant c also has bounded expansion, albeit by a different function.

Similarly, the lexicographic product with a constant-sized graph does not change the expansion of a class arbitrarily.

Proposition 3 (Nešetřil and Ossona de Mendez [73]). *For every graph G, integer $p \geqslant 2$ and half-integer r it is true that*

$$\widetilde{\nabla}_r(G \bullet K_p) \leqslant \max\{2r(p-1)+1, p^2\} \cdot \widetilde{\nabla}_r(G) + p - 1$$

In conclusion, bounded expansion is a robust description of structural sparseness: it is closed under taking shallow minors and the above two 'algebraic' operations. However, bounded expansion classes are not the ultimate definition of sparseness. In order to obtain a dichotonomical view of structural sparseness, Nešetřil and Ossona de Mendez defined *nowhere dense* classes. Their definition mirrors that of bounded expansion classes, only that we now measure the clique number of shallow minors instead of their density.

Definition 5 (Nowhere dense). *A graph class $\mathcal{G}$ is* nowhere dense *if and only if there exists a function f such that for all $r \geqslant 0$, we have $\omega(\mathcal{G} \widetilde{\nabla} r) < f(r)$.*

See the book by Nešetřil and Ossona de Mendez for further definitions [73] and the proof that nowhere dense classes indeed provide a dichotomy of structural sparseness. We will need the following crucial result proved independently by Dvořák and Jiang.

Theorem 1 (Dvořák [29], Jian [46]). *Let $\ell \in \mathbf{N}$ and $\varepsilon > 0$. There exists integers $n_{\ell,\varepsilon}$ and c_ε such that every graph G with $n > n_{\ell,\varepsilon}$ vertices and at least $n^{1+\varepsilon}$ edges contains a c_ε-subdivision of K_ℓ.*

2.2 Random graph models

We usually denote random variables by upper-case letters. Probabilities are denoted by $\mathrm{P}[\bullet]$, expectation, variance, and median by $\mathrm{E}[\bullet]$, $\mathrm{Var}[\bullet]$, $\mathrm{M}[\bullet]$, respectively. If we need to clarify which probability measures we employ, we use subscripts like $\mathrm{P}[\bullet]_M$. For a sequence of random variables $(X_n)_{n\in\mathbf{N}}$ and a random variable Y, recall that (X_n) *converges in distribution* to X if it holds that

$$\forall k \lim_{n\to\infty} \mathrm{P}[X_n \leqslant k] = \mathrm{P}[X \leqslant k].$$

We denote the convergence in distribution with $(X_n) \xrightarrow{d} X$.

A *random graph model* is a sequence of random variables $(G_n)_{n\in\mathbf{N}}$ over n-vertex graphs. For simplicity, we fix $V(G_n) = [n]$. The *parametrisation* of the model is a function $\rho\colon \mathbf{N} \to \mathbf{R}^t$ that creates a tuple of t parameters depending on n which in turn determine the probability distribution of each variable G_n. By $G(n, \rho(n))$ we denote the random variable G_n with the probability distribution prescribed by the model with parameters $\rho(n)$. In order to distinguish random models we will introduce superscripts like G^{CL} or G^{KL}.

For a random graph model $G(n, \rho(n))$ and an integer r the notation $G(n, \rho(n)) \mathbin{\widetilde{\nabla}} r$ denotes a random variable over sets of graphs with at most n vertices whose probability distribution is given by

$$\mathrm{P}\left[G(n, \rho(n)) \mathbin{\widetilde{\nabla}} r = A\right] = \sum_{G:A=G\mathbin{\widetilde{\nabla}} r} \mathrm{P}[G(n, \rho(n)) = G],$$

where A is a set of graphs. With this definition the quantity $\widetilde{\nabla}_\bullet$ is well-defined as a rational-valued random variable. As noted above, we study the properties of random graphs in the limit and hence define the property of having bounded expansion as follows.

Definition 6. *A graph model $G(n, \rho(n))$ has* bounded expansion asymptotically almost surely (a.a.s.) *if there exists a function f such that for all $r \geqslant 0$*

$$\lim_{n\to\infty} \mathrm{P}\left[\widetilde{\nabla}_r(G(n, \rho(n))) < f(r)\right] = 1.$$

It has bounded expansion with high probability (w.h.p.) *if for every $c \geqslant 1$ there exists a function f such that, again for all $r \geqslant 0$,*

$$\mathrm{P}\left[\widetilde{\nabla}_r(G(n, \rho(n))) < f(r)\right] \geqslant 1 - O(n^{-c}).$$

The very same definition is possible to define when a random graph model is nowhere dense.

Definition 7. *A graph model $G(n, \rho(n))$ is* a.a.s. nowhere dense *if there exists a function f such that for all $r \geqslant 0$*

$$\lim_{n\to\infty} \mathrm{P}\left[\omega(G(n, \rho(n) \mathbin{\widetilde{\nabla}} r)) < f(r)\right] = 1.$$

It is nowhere dense w.h.p. *if for every $c \geqslant 1$ there exists a function f such that, again for all $r \geqslant 0$,*

$$\mathrm{P}[\widetilde{\nabla}_r(G(n, \rho(n))) < f(r)] \geqslant 1 - O(n^{-c}).$$

The following notions are needed to prove negative results about random models.

Definition 8. *A graph model $G(n, \rho(n))$ is* a.a.s. somewhere dense *if there exists $r \in \mathbf{N}$ such that for all functions f it holds that*

$$\lim_{n\to\infty} \mathrm{P}\left[\omega(G(n,\rho(n)) \mathbin{\widetilde{\nabla}} r) > f(r)\right] = 1.$$

It is not a.a.s. nowhere dense *if there exists $r \in \mathbf{N}$ such that for all functions f it holds that*

$$\lim_{n\to\infty} \mathrm{P}\left[\omega(G(n,\rho(n)) \mathbin{\widetilde{\nabla}} r) > f(r)\right] > 0.$$

Note that the above definition for random graphs is, in contrast to the definition of structural sparseness for graph classes, not a dichotomy: we can easily build an artificial random graph model where the probability that a dense shallow minor exists converges to some value bounded away from zero and one. This might now just seem like a technicality, but we will see two examples of random graph models developed to replicated certain aspects of complex networks that fall exactly in this category.

The following basic lemma will help to simplify the following proofs by enabling us to work with randomly chosen subgraphs.

Lemma 2. *Let $X_1, \ldots, X_n$ be binary random variables and let $S = \sum_{i=1}^{n} X_i$. Let further $I \in [n]$ be uniformly distributed. Then*

$$\mathrm{P}[S \geqslant 1] \leqslant n \cdot \mathrm{P}[X_I = 1].$$

Proof. By Markov's inequality we have that

$$\mathrm{P}[S \geqslant 1] \leqslant E[S] = \sum_{i=1}^{n} \mathrm{P}[X_i = 1].$$

Observe that

$$\mathrm{P}[X_I = 1] = \sum_{i=1}^{n} \mathrm{P}[I = i] \cdot \mathrm{P}[X_i = 1] = \frac{1}{n} \sum_{i=1}^{n} \mathrm{P}[X_i = 1],$$

and hence

$$\mathrm{P}[S \geqslant 1] \leqslant n \cdot \mathrm{P}[X_I = 1]. \qquad \square$$

We apply this statement to subgraphs in a random graph and obtain the following corollary that lies closer to our application.

Corollary 1. *Let $G(n, \rho(n))$ be a random graph model with parametrisation ρ and Π be a graph property. Then*

$$\mathrm{P}[\exists X \subseteq V(G) \colon G[X] \in \Pi] \leqslant \sum_{k=1}^{n} \binom{n}{k} \mathrm{P}\left[G[Y_k] \in \Pi\right],$$

where Y_k is a k-vertex subset of $V(G)$ chosen uniformly at random.

2.3 Asmyptotic degree distributions

Given a graph G, the *degree sequence* $D(G)$ of G is the sequence $(\deg(v))_{v\in G}$. We say that two graphs G_1 and G_2 have the same degree sequence, if $D(G_1) = \pi(D(G_2))$ for some permutation π. A sequence of n integers $(d_i)_{1\leqslant i\leqslant n}$ is a degree sequence if it can be realized by a graph, i.e. there exists a graph G with $D(G) = (d_i)_{1\leqslant i\leqslant n}$. Such sequences are also called *graphical*.

To define degree distributions, consider a random variable D_n describing the degree of a vertex chosen uniformly at random from an n-vertex graph G. The pmf f_n for D_n is then given by

$$f_n(d) = \frac{b_n(d)}{n} = \frac{1}{n}\sum_{v\in V(G)} [\![\deg(v) = d]\!],$$

where $b_n(d)$ denotes the absolute frequency of degree-d vertices in G.

Definition 9 (Degree distribution)**.** *A n-vertex degree distribution is a random variable D with probability mass function f such that*

1. *$f(d) = 0$ for $d \leqslant 0$ and $d \geqslant n-1$, and*
2. *$nf(d) \in \mathbf{N}_0$ for all $d \in \mathbf{N}$.*

We say that a degree sequence $(d_i)_{1\leqslant i\leqslant n}$ *matches* an n-vertex degree distribution D_n with pmf f_n if for every $1 \leqslant k \leqslant n-1$ it holds that

$$\sum_{i=0}^{n} [\![d_i = k]\!] = nf_n(k).$$

Consequently, a graph G *matches* a degree distribution D_n if its degree sequence does. Since we will consider sequences of random graphs we need to introduce a related notation of sequences of degree distributions.

Definition 10 (Degree distribution sequence, limit, sparse)**.** *A* degree distribution sequence *is an infinite sequence (D_n) of n-vertex degree distributions. A random variable D is the* limit *of (D_n) if $(D_n) \xrightarrow{d} D$. We say that $\mathcal{D}$ is* sparse *if $\mathrm{E}[D] < \infty$ and $(\mathrm{E}[D_n])_{n\in\mathbf{N}} \longrightarrow \mathrm{E}[D]$.*

To motivate the definition of *sparse* sequences, note that for a degree sequence D_n we have that

$$\mathrm{E}[D_n] = \sum_d d f_n(d) = \frac{1}{n}\sum_{d=1}^{n-1} d b_n(d),$$

thus for a graph G with degree distribution D_n it holds that $\mathrm{E}[D_n]$ is exactly its average degree $d_{avg}(G)$. The condition of a degree sequence being sparse will not quite suffice to prove structural sparseness: the situation parallels how graph class with bounded average degree can still harbour dense structures. Consider, for example, the degree distribution sequence of the class consisting of all one-subdivided cliques. While the sequence has constant expectation, the graphs matching it are certainly not structurally sparse. This observation motivates the following stronger condition.

Definition 11 (Tail-bound)**.** *A degree distribution sequence (D_n) with limit D has the function h as its* tail-bound *if there exists a constant $\tau \geqslant 0$ such that for all $d \geqslant \tau$ it holds that*

$$\mathrm{P}[\, D \geqslant d \,] = O\Big(\frac{1}{h(d)} \Big).$$

Referring back to Table 1, all listed function have a tail-bound that is at least quadratic. This is self-evident for power-laws with $\gamma \geqslant 2$; for the other functions we simply note that their second moment exists and hence by Chebyshev's inequality satisfy

$$\mathrm{P}[\, D \geqslant d \,] \leqslant \frac{\mathrm{Var}[D^2]}{d^2 + \mathrm{E}[D^2]} = O\Big(\frac{1}{d^2} \Big).$$

for $d \geqslant \mathrm{E}[D]$.

Observation 1. *Let (D_n) be a sparse degree distribution sequence with limit D. Then $\mathrm{M}[D_n] \longrightarrow \mathrm{M}[D]$.*

3 Graph Models with Bounded Expansion

Often, analytic methods are selected based on their behavior when applied to random graphs that are believed to mimic characteristics of the particular networks under consideration. We will not digress into the arguments for and against this methodology, but simply note that it is a *de facto* part of standard practice at this point in time. Accordingly, we must be able to establish whether or not graphs generated by such models have bounded grad. For more information on random graph models, we refer the readers to the surveys in [69, 70, 79, 81]. In this section, we determine whether several such models have bounded expansion (as by Definition 6.)

We show that the popular configuration model [69, 70], including the version with households exhibiting high clustering [9], and the Chung–Lu

model [21, 22], has bounded expansion w.h.p. in the typical parametric range found in its application to complex networks.

Prior work [74] has shown that the Erdős–Rényi model has bounded expansion a.a.s. Unfortunately, empirical analysis of real-world networks (including friendships/social networks, telephone/communication networks, and biological/neural networks) has shown that typical degree distributions are measurably different from the Poisson distribution exhibited by Erdős–Rényi graphs (see [81] and the references therein.)

We built upon the results of Nešetřil, Ossona de Mendez, and Wood in order to obtain two major results applicable to complex networks. First, the configuration and Chung–Lu model can be seen as an extension of Erdős–Rényi graphs that inherit some of the nice mathematical properties (edge probabilities are somewhat independent in the former and strict independent in the latter) while properly replicating the degree distribution of complex networks by design: Both models allow us to *prescribe* the desired distribution. Typical degree distributions found in real-world graphs are listed in Table 1. Since the prescribed distribution is the one parameter for both models, it is not surprising that the properties of these distributions ultimately determine their structural sparseness. In order to generate sparse graphs, it is necessary for the distributions to have a mean that is independent of n. To generate *structurally* sparse graphs, we show that the distribution in question must have a tail that shrinks at least as fast as a cubic polynomial. This particularly excludes power-law distributions with exponent less than three. However, recent findings have shown that *pure* power law distributions seem to be rare and a precise statistical analysis often favors a power law distribution with an exponential cut-off [24]. Since the latter have in particular tails that are dominated by any polynomial, we conclude that our findings are applicable to the majority of complex networks.

The other major result pertains to what we call the *perturbed bounded- degree model*, which allows the inclusion of an arbitrary (or random) bounded degree graph in addition to probabilistically generated edges (with non- identical probabilities, subject to a uniform bound). We show this model to have bounded expansion with high probability. This in particular strengthens the aforementioned result on the Erdős–Rényi model in terms of speed of convergence. Including a base graph and allowing non-uniform edge probabilities drastically increases the structural variability in the model's output. We will further argue that using graphs with many vertices of unbounded degree as the base graph will necessarily generate structurally dense graphs. Since this model includes as a special case of certain types of stochastic block models, it is relevant to the field of complex network.

Name	Definition $f(d)$	Parameters
Power law	$d^{-\gamma}$	$\gamma > 2$
Power law w/ cutoff	$d^{-\gamma}e^{-\lambda d}$	$\gamma > 2$, $\lambda > 0$
Exponential	$e^{-\lambda d}$	$\lambda > 0$
Stretched exponential	$d^{\beta-1}e^{-\lambda d^{\beta}}$	$\lambda, \beta > 0$
Gaussian	$\exp(-\frac{(d-\mu)^2}{2\sigma^2})$	μ, σ
Log-normal	$d^{-1}\exp(-\frac{(\log d-\mu)^2}{2\sigma^2})$	μ, σ

Table 1: A selection of established functions used to model degree distributions of complex networks, listed without the necessary normalization factors. Here $f(d)$ is the fraction of nodes which have degree d. These functions were taken from an empirical analysis of degree distributions in real-world networks [24].

3.1 The Chung–Lu and configuration model

Given a degree distribution sequence $\mathcal{D} = (D_n)_{n\in\mathbf{N}_0}$ with limit D and an integer n, we are faced with the task to sample graphs uniformly at random from the set

$$\{G \mid G \text{ has degree distribution } D_n\}.$$

Two methods to accomplish this task—with certain caveats—have been put forward: the *configuration model* as described by Bender and Canfield [11] and the model proposed by Chung and Lu [21,22]. The latter is a special case of what has been discussed in the mathematical literature as inhomogeneous random graphs (see, e.g., the work by Bollobás, Janson, and Riordan on the phase transition of such graph [13]).

To sample a graph according to the configuration model, we proceed as follows:

1. Build a degree sequence $(d_i)_{1\leqslant i\leqslant n}$ that matches D_n.
2. Construct a vertex set $V_C = \{v_i^1, \dots, v_i^{d_i}\}_{1\leqslant i\leqslant n}$, i.e. create d_i copies (called *stubs*) for what will be the vertex v^i in the final graph.
3. Generate an auxiliary graph H with vertex set V_C and a random matching as its edge set.
4. Assemble the multi-graph G' with vertex set $\{v_i\}_{1\leqslant i\leqslant n}$ and

   $$|E(v_i, v_j)| = |E_H(\{v_i^1, \dots, v_i^{d_i}\}, \{v_j^1, \dots, v_j^{d_j}\})|,$$

   that is, we connect the vertices v_i and v_j with as many edges as we find between their respective copy-classes in H.

5. Return the graph G derived from G' by removing all parallel edges and loops.

Graphs generate this way will henceforth be denoted by $G^{\mathrm{CF}}(D_n)$. The name *stubs* derives from the following picture of the process: we affix to every vertex a number of half-edges, the stubs, that matches its degree according to the generated degree sequence, and then obtain the multi-graph by randomly wiring the stubs to each other. Instead of generating a matching between the stubs, we can instead choose a random permutation of them and match them up pair-by-pair according to that permutation. This view of the process will be beneficial later on.

It is a priori not clear that this procedure generates graphs with the correct degree sequence. The intermediate multi-graph G' trivially exhibits the degree sequence $(d_i)_{1\leqslant i\leqslant n}$ and thus has the degree distribution D_n. The last step, however, might skew the result by removing parallel edges and loops—we therefore need the probability that G' contains such offending edges to be reasonably low. The conditions under which this is the case have been proved by Molloy and Reed [69] and subsequently improved by Janson, whose result we present here using our own notation.

Theorem 2 (Janson [45])**.** *Let $(D_n)_{n\in\mathbf{N}_0}$ be a degree distribution sequence. Then we have that*

$$\liminf_{n\to\infty} \mathrm{P}[\, G^{CF}(D_n) \text{ is simple}] > 0 \iff \mathrm{E}[D_n^2] = O(\mathrm{E}[D_n]).$$

The original formulation of the theorem's condition is that

$$\sum_{d\geqslant 0} d^2 b_n(d) = O\big(\sum_{d\geqslant 0} d b_n(d)\big).$$

Note that in the case of sparse degree distribution sequences this condition is equivalent to saying that $\mathrm{Var}[D]$ is finite—this already limits the possible degree distributions to those that have a finite variance. Luckily, this is the case for all degree distributions listed in Table 1 with the exceptions of the power law where only for $\gamma > 3$ the variance is finite.

The second method to sample graphs with a prescribed degree distribution, the Chung–Lu model, forgoes the above problems by generating graphs whose *expected* degree distribution matches the given one. Given D_n, it constructs a random graph as follows:

1. Build a degree sequence $(d_i)_{1\leqslant i\leqslant n}$ that matches D_n. We will call d_i the *weight* of the vertex i.

2. Create a graph on n vertices $v_1, \dots v_n$ and connect each pair of vertices v_i, v_j with probability $d_i d_j / m$ where $m = \sum_{k=0}^{n} d_k$.

As network models, both the configuration and the Chung–Lu model suffer from some shortcomings. While they, by design, generate graphs with the correct degree-distribution and small diameter, other statistics found in complex networks are not replicated. In particular, both models have a vanishing clustering-coefficient (see, e.g., Newman's survey [75]).

Since this statistic is critical in e.g. disease propagation research, methods to 'fix' these models have been put forward. A notable example is the *configuration model with household structure* as defined by Ball, Sirl, and Trapman [9]. For this variant, one samples a graph with a prescribed degree sequence and then replaces every vertex by a constant-sized 'household'-graph (for example a clique), distributing the edges incident to a household uniformly to the vertices that comprise it. The resulting graph has provably a constant clustering coefficient.

The main goal of this section will be the proof of the following theorems. To this end, we will first show the behaviour of the Chung–Lu random graphs with respect to $\widetilde{\nabla}_0$ and ω. By an appropriate coupling argument, we can then relate the statistics $\widetilde{\nabla}_r$ and $\omega(\bullet \,\widetilde{\nabla}\, r)$ of different graph models to these base cases. As it turns out, the structural sparseness of the Chung–Lu and the configuration model is entirely determined by the *tail* of the prescribed degree distribution:

Theorem 3. *Let (D_n) be a sparse degree distribution sequence with tail $h(d)$. Both the configuration model $G^{CF}(D_n)$ and the Chung–Lu model $G^{CL}(D_n)$, with high probability,*

- *have bounded expansion for $h(d) = \Omega(d^{3+\varepsilon})$,*
- *are nowhere dense (with unbounded expansion) for $h(d) = \Theta(d^{3+o(1)})$,*
- *and are somewhere dense for $h(d) = O(d^{3-\varepsilon})$.*

Since we can emulate household structures by simply taking the lexicographic product with some constant-size clique and then taking a subgraph, Theorem 3 immediately implies the same for those variants.

Corollary 2. *Let (D_n) be a sparse degree distribution sequence. Then the configuration model $G^{CF}(D_n)$ as well as the Chung–Lu model $G^{CL}(D_n)$ with households have bounded expansion w.h.p. if the sequence has a supercubic tail-bound and nowhere dense w.h.p. if it has a cubic tail-bound.*

In both the Chung–Lu and the configuration model, the first phase consist of assigning weights to vertices according to a degree distribution D_n. This random experiment chooses degrees without replacement, therefore we need to ensure that the important properties of D_n carry over even if a fraction of

the vertices have been uncovered already, that is, we know their weight and hence cannot assume they are randomly distributed.

Since in the following proofs low weights are always preferable, we consider a 'worst-case' variable describing the degree of a vertex after at most n/c other vertices have been assigned the lowest available degrees. Luckily, sparse degree distributions are robust under truncation up to the median. Using this idea, the following lemma shows that we can, up to a point, assume that the vertex degrees are drawn independently according to a modified distribution.

Lemma 3. *Let (D_n) be a sparse degree distribution sequence with limit D and a tail-bound h. Let $\mu_{1/2} = \mathrm{M}[D]$ be the median of D. Then the sequence $(\hat{D}_n)$ defined via*

$$\mathrm{P}[\hat{D}_n = d] = \mathrm{P}[D_n = d \mid D_n \geqslant \mu_{1/2}]$$

is sparse and has tail-bound h.

Proof. Let $\hat{D}$ be defined as the conditioned random variable $D \mid D > \tau$. Then we have that

$$\begin{aligned}\mathrm{E}[\hat{D}] &= \sum_{d>0} d\,\mathrm{P}[D = d \mid D \geqslant \mu_{1/2}] = \sum_{d \geqslant \mu_{1/2}} d \cdot \frac{\mathrm{P}[D = d]}{\mathrm{P}[D \geqslant \mu_{1/2}]} \\ &= \frac{\mathrm{E}[D]}{\mathrm{P}[D \geqslant \mu_{1/2}]} \leqslant \frac{\mathrm{E}[D]}{2}.\end{aligned}$$

Hence $\mathrm{E}[\hat{D}]$ is finite and $(\hat{D}_n)$ is sparse.

Let τ be the constant such that for $d \geqslant \tau$ it holds that $\mathrm{P}[D \geqslant d] = O(1/h(d))$. To see that the same holds for $(\hat{D}_n)$, note that

$$\mathrm{P}[\hat{D} \geqslant d] = \mathrm{P}[D \geqslant d \mid D \geqslant \mu_{1/2}] = \frac{\mathrm{P}[D \geqslant \max\{d, \mu_{1/2}\}]}{\mathrm{P}[D \geqslant \mu_{1/2}]}.$$

Hence for $\hat{\tau} \geqslant \max\{\tau, \mu_{1/2}\}$ for all $d \geqslant \hat{\tau}$ the bound

$$\mathrm{P}[\hat{D} \geqslant d] = O\big(\frac{1}{h(d)}\big)$$

holds, as claimed. □

In the remainder of this section, we work towards a proof of Theorem 3. Ultimately, our approach couples the density of shallow minors for graphs with degree distribution D to the density of subgraphs of a random graph with distribution ηD, where η is an appropriate scaling factor. This factor ultimately traces back to the value of the sum $\sum_{d=1}^{\Delta} \frac{1}{d^{\gamma-2}}$ which is a constant for $\gamma > 3$, roughly $\log \Delta$ for $\gamma = 3$ and around $\Delta^{3-\gamma}$ for $\gamma < 3$. We will make use of the following asymptotic bound for such sums.

Lemma 4. *For $\alpha \geqslant 1$, an integer $\delta > 0$, and an unbounded positive function $\Delta(n)$ it holds that*

$$\sum_{k=\delta}^{\Delta} \frac{1}{k^{\alpha+1}} \sim \frac{1}{\alpha\delta^{\alpha}}.$$

Proof. The function $k^{-(\alpha+1)}$ is monotonically decreasing on $(0, \infty)$ and therefore it is true that

$$\int_{\delta}^{\Delta} \frac{1}{k^{\alpha+1}} \leqslant \sum_{k=\delta}^{\Delta} \frac{1}{k^{\alpha+1}} \leqslant \frac{1}{\Delta^{\alpha+1}} + \int_{\delta}^{\Delta} \frac{1}{k^{\alpha+1}}.$$

The integrals simply evaluate to

$$\int_{\delta}^{\Delta} \frac{1}{k^{\alpha+1}} = \Big[-\frac{1}{\alpha k^{\alpha}} \Big]_{\delta}^{\Delta} = \frac{1}{\alpha\delta^{\alpha}} - \frac{1}{\alpha\Delta^{\alpha}} \sim \frac{1}{\alpha\delta^{\alpha}}.$$

Since $\Delta^{-(\alpha+1)}$ converges to zero, the sum is upper- and lower-bounded by functions that are on the order of $1/\alpha\delta^{\alpha}$ and the claim follows. □

3.2 The supercubic regime

We begin our investigation by considering degree distributions whose tail is bounded by a supercubic function. As discussed earlier, most real-world degree distributions seem to fall into this pattern since they exhibit a sharp *exponential* cut-off.

Lemma 5. *Consider vertices $\{v_i\}_{i\in[k]}$ with associated weights $\{d_i\}_{d\in[\Delta]}$, for some $\Delta \in \mathbf{N}$. Let G be a random graph on these vertices where each edge v_iv_j is independently present with probability $\leqslant \beta d_i d_j/n$. Then*

$$\mathrm{P}[\,\|G\| \geqslant \xi k\,] \leqslant \Big(\frac{e\beta D^2}{2n\xi k e^{D^2/2n}}\Big)^{\xi k}$$

if $D := \sum_i d_i$ satisfies $\beta D^2 \leqslant 2n\xi k$.

Proof. We associate a random variable X_{ij} with every edge v_iv_j. The expected number of edges is then

$$E[\sum_{ij} X_{ij}\,] = \sum_{ij} E[\,X_{ij}\,] \leqslant \sum_{ij} \frac{\beta d_i d_j}{n} = \frac{\beta}{2n}\big(\sum_i d_i\big)^2 = \frac{\beta D^2}{2n}.$$

We apply the Chernoff-bound

$$\mathrm{P}\Big[\sum_{ij} X_{ij} \geqslant (1+\delta)\frac{\beta D^2}{2n}\Big] \leqslant \Big(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\Big)^{\beta D^2/2n}$$

choosing $\delta = \frac{2n\xi k}{\beta D^2} - 1$ and obtain

$$\mathrm{P}\Big[\sum_{ij} X_{ij} \geqslant \xi k\Big] \leqslant \frac{e^{\xi k - \beta D^2/2n\xi k}}{(\frac{2n\xi k}{\beta D^2})^{\xi k}} = \Big(\frac{e\beta D^2}{2n\xi k e^{D^2/2n}}\Big)^{\xi k},$$

as claimed. □

The first important puzzle piece is that the assumed tail-bound ensures the presence of only few high-weight vertices:

Lemma 6. *Let (D_n) be a sparse degree distribution sequence with limit D with a polynomial tail-bound $h(d) = d^{\alpha+1}, \alpha > 1$. Then for every graph G that realizes D_n and any set $X \subseteq V(G)$ of k vertices, the sum-of-degrees of X is at most*

$$\sum_{v \in X} \deg_G(v) \leqslant \sqrt{2} n^{1/\alpha} k^{\alpha - 1/\alpha}.$$

Proof. Let $\Delta = h^{-1}(n) = n^{1/\alpha+1}$ be the maximal realizable degree. To maximize the sum-of-degrees we obviously need to pick the k vertices of highest degree. Hence we want the largest threshold δ such that

$$\sum_{d=\delta}^{\Delta} \frac{n}{d^{\alpha+1}} \geqslant k.$$

Using the estimate provided by Lemma 4, we see that

$$\sum_{d=\delta}^{\Delta} \frac{n}{d^{\alpha+1}} \sim \frac{n}{\alpha\delta^{\alpha}} \quad \Longrightarrow \quad \delta \sim \big(\frac{n}{\alpha k}\big)^{1/\alpha}.$$

Given δ, the maximal possible sum-of-degrees is at most

$$\begin{aligned}\sum_{d=\delta}^{\Delta} \frac{n}{d^{\alpha+1}} \cdot d = \sum_{d=\delta}^{\Delta} \frac{n}{d^{\alpha}} &\sim \frac{n}{(\alpha-1)\delta^{\alpha-1}} \\ &= \frac{\alpha^{(\alpha-1)/\alpha}}{\alpha - 1} \cdot \frac{n k^{(\alpha-1)/\alpha}}{n^{(\alpha-1)/\alpha}} \leqslant \sqrt{2} n^{1/\alpha} k^{\alpha - 1/\alpha},\end{aligned}$$

as claimed. □

Using this property, we can now show that Chung–Lu graphs generated with degree distributions that have a supercubic tail will not contain dense subgraphs. Afterwards we will related the general case back to this basic case by coupling.

Lemma 7. *Let (D_n) be a sparse degree distribution sequence with limit D with a polynomial tail-bound $h(d) = d^{\alpha+1}, \alpha > 2$. Then there exists a function f such that for every $\xi \geqslant 4ae^2$, every $c \geqslant 4e$ and every $n \geqslant f(\xi, \alpha)$ it holds that*

$$\mathrm{P}[\exists H \subseteq G^{CL}(D_n)) : |H| \leqslant n/c \text{ and } \nabla_0(H) \geqslant \xi] \leqslant \frac{1}{n^{\xi}}.$$

Proof. Using Lemma 2, we can bound the probability that a dense subgraph on k vertices exists by considering the probability that k randomly chosen vertices form a dense subgraph. Taking the union bound of all possible k (note that we need at least $2\xi + 1$ vertices for a subgraph of density ξ, we simplify this lower bound to 2ξ), the probability of the aforementioned event is at most

$$\sum_{k=2\xi}^{n/c} \binom{n}{k} \mathrm{P}[\|G^{\mathrm{CL}}(D_n)[X_k])\| \geqslant k\xi],$$

where X_k is a set of k vertices chosen uniformly at random.

By applying Lemma 5 to the random subgraph $G^{\mathrm{CL}}(D_n)[X_k]$, we can bound the above probability by

$$\begin{aligned}
&\sum_{k=2\xi}^{n/c} \binom{n}{k} \sum_{d=k}^{k\Delta} \left(\frac{ecd^2}{2n\xi k e^{d^2/2n}}\right)^{\xi k} \mathrm{P}[D_k = d] \\
&\leqslant \sum_{k=2\xi}^{n/c} \frac{n^k e^k}{k^k} \left(\frac{ec}{2n\xi k}\right)^{\xi k} \sum_{d=k}^{k\Delta} \frac{d^{2\xi k}}{e^{\xi k d^2/2n}} \mathrm{P}[D_k = d],
\end{aligned}$$

where D_k denotes the degree-sum of k vertices chosen from D_n.

By Lemma 6 we know that for a degree-distribution with tail-bound $h(d) = d^{\alpha+1}$ the degree-sum of k vertices cannot exceed $\hat{D}_k := \sqrt{2}n^{1/\alpha}k^{\alpha-1/\alpha}$. Since

$$\sum_{d=k}^{\hat{D}_k} d^{2\xi k} \leqslant \hat{D}_k^{2\xi k+1} = \left(\sqrt{2}n^{1/\alpha}k^{\alpha-1/\alpha}\right)^{2\xi k+1},$$

we can upper-bound every term of the above sum by

$$\begin{aligned}
&\frac{n^k e^k}{k^k} \left(\frac{ec}{2n\xi k}\right)^{\xi k} \left(\sqrt{2}n^{1/\alpha}k^{\alpha-1/\alpha}\right)^{2\xi k+1} \\
&\leqslant \frac{(2ec)^{(\xi+1)k}}{(2\xi)^{\xi k}} \cdot \frac{n^{\frac{1}{\alpha}2\xi k+1} k^{\frac{\alpha-1}{\alpha}(2\xi k+1)}}{n^{(\xi-1)k} k^{(\xi+1)k}}.
\end{aligned}$$

Choosing $\xi >$ large enough, the first factor can be decreased below one (in particular, this is true for $\xi \geqslant 4e^2$). Hence, we focus on the second factor:

$$\frac{n^{\frac{1}{\alpha}(2\xi k+1)} k^{\frac{\alpha-1}{\alpha}(2\xi k+1)}}{n^{(\xi-1)k}} = \frac{k^{(\frac{\alpha-1}{\alpha}2\xi-\xi-1)k+\frac{\alpha-1}{\alpha}}}{n^{(\xi-\frac{1}{\alpha}2\xi-1)k-\frac{1}{\alpha}}} = \frac{k^{(\frac{\alpha-2}{\alpha}\xi-1)k+\frac{\alpha-1}{\alpha}}}{n^{(\frac{\alpha-2}{\alpha}\xi-1)k-\frac{1}{\alpha}}}.$$

Therefore the probability that a subgraph of density at least ξ exists is at most

$$\sum_{k=2\xi}^{n/c} \frac{k^{(\frac{\alpha-2}{\alpha}\xi-1)k+\frac{\alpha-1}{\alpha}}}{n^{(\frac{\alpha-2}{\alpha}\xi-1)k-\frac{1}{\alpha}}}.$$

Consider the ratio of two consecutive terms: we see that

$$\begin{aligned}
&\frac{(k-1)^{(\frac{\alpha-2}{\alpha}\xi-1)(k-1)+\frac{\alpha-1}{\alpha}}}{n^{(\frac{\alpha-2}{\alpha}\xi-1)(k-1)-\frac{1}{\alpha}}} \cdot \frac{n^{(\frac{\alpha-2}{\alpha}\xi-1)k-\frac{1}{\alpha}}}{k^{(\frac{\alpha-2}{\alpha}\xi-1)k+\frac{\alpha-1}{\alpha}}} \\
&= \left(\frac{n}{k}\right)^{\frac{\alpha-2}{\alpha}\xi-1} \cdot \left(1-\frac{1}{k}\right)^{(\frac{\alpha-2}{\alpha}\xi-1)(k-1)+\frac{\alpha-1}{\alpha}} \\
&\geqslant \left(\frac{n}{k}\right)^{\frac{\alpha-2}{\alpha}\xi-1} \left(\frac{1}{e}\right)^{\frac{\alpha-2}{\alpha}\xi-1} \left(\frac{1}{2}\right)^{\frac{\alpha-1}{\alpha}}.
\end{aligned}$$

Since k is at most n/c, this ratio is at least

$$\frac{c^{\frac{\alpha-2}{\alpha}\xi-1}}{e^{\frac{\alpha-2}{\alpha}\xi-1}2^{\frac{\alpha-1}{\alpha}}} \geqslant \left(\frac{c}{2e}\right)^{\frac{\alpha-2}{\alpha}\xi-1},$$

where the last simplification holds for $\xi \geqslant 4e^2 > 2\frac{\alpha-1}{\alpha-2}$. We conclude that for $c \geqslant 4e$ the whole ratio is at least 2; and accordingly the bound

$$\sum_{k=2\xi}^{n/c} \frac{k^{(\frac{\alpha-2}{\alpha}\xi-1)k+\frac{\alpha-1}{\alpha}}}{n^{(\frac{\alpha-2}{\alpha}\xi-1)k-\frac{1}{\alpha}}} \leqslant 2\frac{(2\xi)^{(\frac{\alpha-2}{\alpha}\xi-1)2\xi+\frac{\alpha-1}{\alpha}}}{n^{(\frac{\alpha-2}{\alpha}\xi-1)2\xi-\frac{1}{\alpha}}} \leqslant \frac{1}{n^{\xi}}.$$

holds, where the last inequality follows for large enough n (depending only on α, ξ). □

We pair the above lemma and Proposition 2 with a coupling argument to arrive at the following result. For simplicity, we define $\mathcal{E}^X_{\xi,r}$ as the event that a random graph contains an r-shallow topological minor of density at least ξ whose embedding ϕ_V maps into X.

Lemma 8. *Let $G^R(n)$ be a random graph model with the following property: for every r there exists a sparse degree distribution (D_n) with tail-bound $h(d) = d^{\alpha+1}, \alpha > 2$ such that for every $\xi \geqslant 4ae^2$ it holds that*

$$\mathrm{P}[\mathcal{E}^X_{\xi,r}] \leqslant \mathrm{P}[\nabla_0(G^{CL}(D_n)[X]) \geqslant \xi],$$

where X is a random set of at most $n/4e(r\xi+1)$ vertices. Then $G^R(n)$ has bounded expansion with high probability.

Proof. Note that if the event $\mathcal{E}^X_{\xi,r}$ occurs, it already occurs in a subgraph of size $|X|+r\xi|X|$. Therefore the maximal size of X that needs to be considered in order to apply Proposition 2 is

$$|X| + r\xi|X| \leqslant \frac{n}{4e} \iff |X| \leqslant \frac{n}{4e(r\xi+1)}.$$

Exchanging the probability $\mathrm{P}[\mathcal{E}^X_{\xi,r}]$ by $\mathrm{P}[\widetilde{\nabla}_r(G^{\mathrm{R}}(n)[X]) \geqslant \xi]$ in the proof of Lemma 7 immediately shows that

$$\sum_{k=2\xi}^{n/4e(r\xi+1)} \binom{n}{k} \mathrm{P}[\mathcal{E}^X_{\xi,r}] \leqslant \frac{1}{n^\xi}$$

for suitably large n. By Lemma 2, therefore the probability that any set of at most $n/4e$ vertices form the nails of a dense r-shallow minor is at most $n^{-\xi}$. Accordingly, setting $f_{\widetilde{\nabla}} = \xi$ and $f_H = 1/4e$, the second condition of Proposition 2 holds with probability at least $1 - n^{-\xi}$. It is left to show that functions $f_{\mathrm{thresh}}, f_{\mathrm{deg}}$ exist.

Claim. *Let (f_n) be the pmfs and D the limit of (D_n). Then every graph matching (D_n) satisfies Condition 1 of Proposition 2.*

Recall that Condition 1 states that there exist functions $f_{\mathrm{thresh}}, f_{\mathrm{deg}}$ such that for all ε we either have $|G| \leqslant f_{\mathrm{thresh}}(\varepsilon)$ or it holds that

$$|\{v \in V(G)\colon \deg(v) \geqslant f_{\mathrm{deg}}(\varepsilon)\}| \leqslant \varepsilon \cdot |G|.$$

This translates to D_n as follows: for every $\varepsilon > 0$ there exists an integer $0 \leqslant d \leqslant n-1$ such that

$$n \sum_{k=d}^{n-1} f_n(k) \leqslant \varepsilon n \iff \sum_{k=d}^{n-1} f_n(k) \leqslant \varepsilon.$$

We apply Markov's inequality to find that

$$\sum_{k=d}^{n-1} f_n(k) = \mathrm{P}[D_n \geqslant d] \leqslant \frac{\mathrm{E}[D_n]}{d}.$$

Since $\mathrm{E}[D_n] \longrightarrow \mathrm{E}[D]$ and $\mathrm{E}[D]$ is finite, the right hand side can be made small enough by choosing $f_{\mathrm{deg}}(\varepsilon) = d = \mathrm{E}[D_n]/\varepsilon$ and n large enough. This proves the existence of appropriate functions f_{thresh} and f_{deg} and the claim.

Hence, we conclude that Proposition 2 is applicable to $G^{\mathrm{R}}(n)$ with probability at least $(1 - n^{-\xi})$ and the claim follows. □

Having shown the supercubic coupling lemma, we proceed to the next range of distributions.

3.3 The cubic regime

As before, we will work towards a coupling argument which translates probabilities for events on shallow topological minors to events on subgraphs. Starting at the bottom, we begin our investigation by proving an analogue of Lemma 7 for degree distributions that are scaled by a factor of $\log^{\Theta(1)} n$.

Lemma 9. *Let (D_n) be a sparse degree distribution sequence with limit D and tail-bound $h(d) = \Theta(d^{3+o(1)})$. Then for $\xi > 9$ it holds that*

$$\mathrm{P}[\omega(G^{CL}(\log^{\Theta(1)} nD_n)) \geqslant \xi] = O(n^{-\xi}).$$

Proof. Let $G = G^{\mathrm{CL}}(\log^{\Theta(1)} nD_n)$. Using Lemma 2, we can bound the probability that a clique on k vertices exists by considering the probability that k randomly chosen vertices form a clique. Taking the union bound of all possible k, the probability of the aforementioned event is at most

$$\sum_{k=\xi}^{n} \binom{n}{k} \mathrm{P}[\omega(G[X_k]) \geqslant k],$$

where X_k is a set of k vertices chosen uniformly at random. In the following, we will talk about the vertex weights assigned in the creation of the Chung–Lu random graph G by factoring out the common $\log^{\Theta(1)} n$-factor: a vertex v has weight $d_v \log^{\Theta(1)} n$.

An edge uv exists in G with probability $d_u d_v \log^{\Theta(1)} n/\mu n$, where again $\mu = E[D_n]$. For k vertices of weights $d_1, \dots, d_k$, the probability that they form a complete graph is therefore at most

$$\prod_{i<j} \frac{d_i d_j \log^{\Theta(1)} n}{\mu n} = \Big(\frac{\log^{\Theta(1)} n}{\mu n}\Big)^{\frac{k(k-1)}{2}} \Big(\prod_i d_i\Big)^{k-1}.$$

Note that the tail-bound H of $\log^{\Theta(1)} nD_n$ is

$$H(d) = h\Big(\frac{d}{\log^{\Theta(1)} n}\Big) = \Big(\frac{d}{\log^{\Theta(1)} n}\Big)^{3+o(1)},$$

hence the maximal degree is $\Delta = H^{-1}(n) = n^{1/(3+o(1)))} \log^{\Theta(1)} n$.

Taking the union bound of all possible products D of weight k, we obtain the bound

$$\begin{aligned}\sum_{k=\xi}^{n} \binom{n}{k} \mathrm{P}[\omega(G[X_k]) \geqslant k] &\leqslant \sum_{k=\xi}^{n} \binom{n}{k} \Big(\frac{\log^{\Theta(1)} n}{\mu n}\Big)^{\frac{k(k-1)}{2}} \sum_{D=k}^{\Delta^k} D^{k-1} \\ &\leqslant \sum_{k=\xi}^{n} \frac{n^k e^k}{k^k} \Big(\frac{\log^{\Theta(1)} n}{\mu n}\Big)^{\frac{k(k-1)}{2}} \Delta^{k^2}.\end{aligned}$$

Is is easy to check that this sum is supergeometric, hence we can estimate its value by twice its largest term. This results in

$$\sum_{k=\xi}^{n} \binom{n}{k} \mathrm{P}[\omega(G[X_k]) \geqslant k] \leqslant 2\frac{n^\xi e^\xi}{\xi^\xi}\Big(\frac{\log^{\Theta(1)} n}{\mu n}\Big)^{\frac{\xi(\xi-1)}{2}} \Delta^{\xi^2}$$
$$= \frac{2e^\xi}{(\xi\mu)^\xi}\frac{(\log^{\Theta(1)} n)^{\frac{3}{2}\xi^2-\frac{1}{2}\xi}}{n^{(1/2-1/(3+o(1)))\xi^2-\frac{3}{2}\xi}}.$$

Since this last term is, for $\xi > 9$ and large enough n, lies in $O(n^{-\xi})$, the claim follows. □

Let in the following $\mathcal{K}_r^X$ denote the event that the vertices of X form the nails of an r-subdivision of a complete graph. Then we can rephrase the above lemma as follows:

Corollary 3. *Let $G^R(n)$ be a random graph model with the following property: for every r there exists a sparse degree distribution (D_n) with tail-bound $h(d) = \Theta(d^{3+o(1)})$ and $\xi > 9$ such that*

$$\mathrm{P}[\mathcal{K}_r^X] \leqslant \mathrm{P}[G^{CL}(\log^{\Theta(1)} nD_n)[X] \simeq K_{|X|}],$$

where X is a random set of at most ξ vertices. Then $G^R(n)$ is nowhere dense with high probability.

The above lemma will later provide the positive statement, namely, that large shallow clique-minors are improbable due to a coupling to the occurrence to large cliques in the above setting. However, Theorem 3 also states that the cubic degree distribution do *not* result in graphs with bounded expansion. The following lemma provides us with the necessary negative statement.

Lemma 10. *Let (D_n) be a sparse degree distribution sequence with tail-bound $h(d) = \Theta(d^{3+o(1)})$. Then*

$$\widetilde{\nabla}_{1/2}(G^{CL}(D_n)) = \Omega(\log n)$$

with high probability.

Proof. Consider the $\log n$ vertices V_h of highest weight in $G^{\mathrm{CL}}(D_n)$. Then the minimum weight δ of these vertices can be estimated by seeing that (using again the bound provided by Lemma 4)

$$N = \sum_{d=\delta}^{\Delta} \frac{n}{d^{3+o(1)}} \sim \frac{n}{(2+o(1))\delta^{2+o(1)}}$$
$$\Longrightarrow \quad \delta \sim \Big(\frac{n}{(2+o(1))\log n}\Big)^{1/(2+o(1))}.$$

We will show that there exists a dense $1/2$-shallow topological minor with V_h as its nails. Note that at most $\log^2 n$ vertices can participate as subdivision-vertices in this minor. By assuming that only the $n-(\log n+\log^2 n)$ vertices of lowest weight can be used as branch vertices, we can simulate the density of the minor in a simpler random graph: we connect vertices $s, t \in V_h$ independently at random with probability p such that

$$p \leqslant \mathrm{P}[\exists w \in V(G) \setminus V_h : sw, st \in G^{\mathrm{CL}}(D_n)].$$

Clearly, the surrogate graph's density is (stochastically) smaller than the minor's density.

First, let us estimate the maximal available vertex weight according to our worst-case assumptions: we need a value η that satisfies

$$\begin{aligned} &\frac{n}{(2+o(1))\eta^{2+o(1)}} \geqslant \log n + \log^2 n \\ \implies \quad &\eta \leqslant \Big(\frac{n}{(2+o(1))(\log n + \log^2 n)}\Big)^{1/(2+o(1))}. \end{aligned}$$

For simplicity, we choose $\eta = (n/4log^2 n)^{1/(2+o(1))}$ which satisfies the above constraint for large enough n. Now the probability that two nails are both connected to a branch vertex is at least

$$\begin{aligned} p &:= \sum_{d=1}^{\eta} \frac{n}{d^{3+o(1)}} \cdot \frac{\delta^2 d^2}{\mu^2 n^2} = \frac{\delta^2}{n} \sum_{d=1}^{\eta} \frac{1}{d^{1+o(1)}} \\ &\sim \frac{\delta^2 \log \eta}{\mu^2 n} \geqslant \frac{\log n - \log(4\log^2 n)}{10\mu^2 n^{o(1)} \log^{1+o(1)} n}. \end{aligned}$$

For large enough n, we therefore have that p is bounded from below by a constant. Hence the expected number of edges in the minor is $\Omega(\log^2 n)$ and its density $\Omega(\log n)$. Since these surrogate edges are independent, their number and hence the lower-bound for the minor's density are sharply concentrated around the mean: the probability of falling below a factor of $(1-t)\log n$ density is only $O(n^{-t^2/2})$. We conclude that the minor has density $\Omega(\log n)$ with high probability. □

Having fully characterised the near-cubic regime, we proceed to the final range; degree distributions with a subcubic tail.

3.4 The subcubic regime

As it turns out, the 'phase transition' from the nowhere-dense regime at d^3 is very sharp: a degree distribution with a tail lower-bounded by $d^{3-\varepsilon}$ for

any $\varepsilon > 0$ will result in the presence of shallow dense clique minors of arbitrary size. Since we can leverage the powerful Theorem 1, the proof is quite straightforward.

Lemma 11. *Let (D_n) be a sparse degree distribution sequence with a tail lower-bounded by $h(d) = O(d^{3-\varepsilon})$ for some $\varepsilon > 0$. Then $G^{CL}(D_n)$ is somewhere dense with high probability.*

Proof. Consider the N vertices of highest degree, where N will be fixed later. Let $\gamma = 3 - \varepsilon$. Then the minimum weight δ of these vertices can be estimated by seeing that (again applying Lemma 4)

$$N = \sum_{d=\delta}^{\Delta} \frac{n}{d^\gamma} \sim \frac{n}{(\gamma-1)\delta^{\gamma-1}} \quad \Longrightarrow \quad \delta \sim \Big(\frac{n}{(\gamma-1)N}\Big)^{1/(\gamma-1)}.$$

Hence the expected number of edges within these N vertices is at least

$$N^2 \frac{1}{\mu n}\Big(\frac{n}{(\gamma-1)N}\Big)^{2/(\gamma-1)} = N^{\frac{2(\gamma-2)}{\gamma-1}} n^{\frac{3-\gamma}{\gamma-1}} \mu^{\frac{2}{\gamma-1}} (\gamma-1)^{-\frac{2}{\gamma-1}}.$$

Our goal is to show that for a suitable choice of N depending on n, the number of expected edges is $N^{1+\varepsilon'}$ for some $\varepsilon' > 0$. This condition yields:

$$\begin{aligned} & N^{\frac{2(\gamma-2)}{\gamma-1}} n^{\frac{3-\gamma}{\gamma-1}} \mu^{\frac{2}{\gamma-1}} (\gamma-1)^{-\frac{2}{\gamma-1}} \geqslant N^{1+\varepsilon'} \\ \iff & n^{\frac{3-\gamma}{\gamma-1}} \mu^{\frac{2}{\gamma-1}} (\gamma-1)^{-\frac{2}{\gamma-1}} \geqslant N^{\frac{3-\gamma}{\gamma-1}+\varepsilon'}. \end{aligned}$$

Choosing, for example, $\varepsilon' = (3-\gamma)/(\gamma-1)$ we can let $N = O(\sqrt{n})$. Now by Theorem 1, this implies that for every p the graph $G^{\mathrm{CL}}(D_n)$ contains a $c_{\varepsilon'}$-subdivision of K_p if we find at least the expected number of edges between the $O(\sqrt{n})$ vertices of highest degree. By the Chernoff bound, this does *not* happen only with probability

$$\mathrm{P}[E_N \leqslant (1-t)n^{\varepsilon'/2}] \leqslant \exp\Big(-\frac{n^{\varepsilon'/2}t^2}{2}\Big) = \exp\Big(-\frac{n^{\varepsilon/(4-2\varepsilon))}t^2}{2}\Big),$$

hence the claim follows. □

3.5 The proof of Theorem 3

We begin with the proof for the Chung–Lu model since the coupling arguments are straightforward. Afterwards, we will show how they can be adapted to extend the proof to the configuration model. The important puzzle piece for the coupling is expressed in the following lemma: the probability that a short path exists in $G^{\mathrm{CL}}(D_n)$ crucially depends on the weight of its endpoints and the shape of the degree distribution's tail.

Lemma 12. *Let (D_n) be a sparse degree-distribution whose tail is bounded by h for degrees above τ. Let s, t be vertices in $G^{CL}(D_n)$. Then for every integer r it holds that*

$$\mathrm{P}[\exists P_{st} \subseteq G^{CL}(D_n), |P_{st}| = r \mid d_s, d_t, F] = \frac{d_s d_t}{n} O(E[D_n^2]^{r-1})$$

and this bound still holds if up to $n/2$ weights have been uncovered.

Proof. Consider the probability that a fixed path $P_{st} := s, v_1, \ldots, v_{r-1}, t$ is realized in $G := G^{\mathrm{CL}}(D_n)$ if the assigned vertex weights are, respectively, $d_s, d_1, \ldots, d_{r-1}, d_t$: this probability is given by

$$\mathrm{P}[P_{st} \subseteq G \mid d_s, d_{v_1}, \ldots, d_{v_r-1}, d_t] = \frac{d_s d_t \prod_i d_i^2}{\mu^r n^r},$$

where $\mu = \mathrm{E}[D_n]$ is independent of n.

Since only a fraction of the vertex weights have been uncovered, instead of choosing weights without replacement, we can use $\hat{D}_n$ from Lemma 3 to sample weights independently. Accordingly, we let the variables $\hat{D}_{1,n}, \ldots, \hat{D}_{r-1,n}$ be independent copies of D_n.

Recall that $(\hat{D}_n)$ has the same tail-bound h as (D_n). Note now that

$$\mathrm{E}[D_n^2] = \sum_{d=1}^{\tau-1} \mathrm{P}[D_n = d] \cdot d^2 + \sum_{d=\tau}^{\Delta} \frac{d^2}{h(d)} = \Theta(\mathrm{E}[\hat{D}_n^2]).$$

Since the $D_{i,n}$ are independent, we can estimate the union-bound over all weights

$$\begin{aligned}
\sum_{d_1,\ldots,d_{r-1}} \prod_i d_i^2 \cdot \mathrm{P}[\hat{D}_{i,n} = d_i] &= \prod_{d_1,\ldots,d_{r-1}} \sum_i d_i^2 \cdot \mathrm{P}[\hat{D}_{i,n} = d_i] \\
&= \prod_{d_1,\ldots,d_{r-1}} E[\hat{D}_{i,n}^2] \\
&= \Theta(E[D_n^2]^{r-1}).
\end{aligned}$$

We arrive at the upper bound

$$\mathrm{P}[P_{st} \subseteq G \mid d_s, d_t, F] = \frac{d_s d_t}{\mu^r n^r} \Theta(E[D_n^2]^{r-1}).$$

Since a path of length r has $r-1$ internal vertices, the probability that *some* s-t-paths of length r exists is (implicitly using Lemma 2) then bounded by

$$\mathrm{P}[\exists P_{st} \subseteq G \mid d_s, d_t] = \frac{d_s d_t}{n} O(E[D_n^2]^{r-1}),$$

as claimed. □

We finally have all the ingredients for the main proof.

Proof of Theorem 3 for $G^{CL}(D_n)$. First consider a sparse degree distribution sequence (D_n) with tail-bound $h(d) = \Omega(d^{3+\varepsilon})$ for some $\varepsilon > 0$. By Lemma 12, the probability of an s-t-path of length r existing in $G := G^{\mathrm{CL}}(D_n)$ is

$$\begin{aligned}\mathrm{P}[\exists P_{st} \subseteq G] &\leqslant \mathrm{P}[st \in G^{\mathrm{CL}}(\Theta\big(\sqrt{E[D_n^2]^{r-1}}\big)D_n)] \\ &\leqslant \mathrm{P}[st \in G^{\mathrm{CL}}(\Theta(D_n))]\end{aligned}$$

and this relation is still true if conditioned by the knowledge of up to $n/4e$ vertex-weights. Hence for any random set X of at most $n/4e$ vertices we have that

$$\mathrm{P}[\mathcal{E}^X_{\xi,r}]_{G^{\mathrm{CL}}(D_n)} \leqslant \mathrm{P}[\nabla_0(G^{\mathrm{CL}}(\sqrt{c_r}D_n)[X])) \geqslant \xi].$$

Thus by Lemma 8 and the fact that $\Theta(D_n)$ is sparse and has the same tail-bound as D_n, the model $G^{\mathrm{CL}}(D_n)$ has bounded expansion with high probability.

Next, assume (D_n) has a tail $h(d) = \Theta(d^{3+o(1)})$. By Lemma 12, the probability of an s-t-path of length r existing in $G := G^{\mathrm{CL}}(D_n)$ is

$$\begin{aligned}\mathrm{P}[\exists P_{st} \subseteq G] &\leqslant \mathrm{P}[st \in G^{\mathrm{CL}}(\Theta\big(\sqrt{E[D_n^2]^{r-1}}\big)D_n)] \\ &\leqslant \mathrm{P}[st \in G^{\mathrm{CL}}(\Theta(\log^{\Theta(1)} nD_n))]\end{aligned}$$

and this relation is still true if conditioned by the knowledge of up to $n/4e$ vertex-weights. Let again $\mathcal{K}^X_r$ denote the event that an r-subdivision of a complete subgraph with nails X exists in G. Since the graph G is sparse with high probability, we focus on the case $r \geqslant 1$. Now for any random set X of at most $\sqrt{n/4er}$ vertices we have that

$$\mathrm{P}[\mathcal{K}^X_r]_{G^{\mathrm{CL}}(D_n)} \leqslant \mathrm{P}[G^{\mathrm{CL}}(\log^{\Theta(1)} D_n)[X]) \simeq K_{|X|}].$$

and thus by Corollary 3 it follows that $G^{\mathrm{CL}}(D_n)$ is nowhere dense with high probability. By Lemma 10, we further have that already the measure $\widetilde{\nabla}_{1/2}(G)$ grows at a rate of at least $\Omega(\log n)$, hence $G^{\mathrm{CL}}(D_n)$ has unbounded expansion.

Finally, assume (D_n) has a tail-bound $h(d) = O(d^{3-\varepsilon})$ for some $\varepsilon > 0$. By Lemma 11 we already have that $G^{\mathrm{CL}}(D_n)$ is somewhere dense with high probability. □

This proof can be extended to the configuration model, the main difficulty here is that edges are not sampled independently of each other. We first prove a variant of Lemma 12. The bound proved here crucially depends on the number of *unmatched* stubs: recall that we can, instead of matching up stubs by choosing a random matching, match them up pair-by-pair. From this perspective, we can stop the process at any point and express the probabilities at this stage in terms of the remaining number of stubs.

Lemma 13. *Let (D_n) be a sparse degree-distribution whose tail is bounded by h for degrees above τ. Let s, t be vertices in $G^{CF}(D_n)$. Then for every $r \in \mathbf{N}$ it holds that*

$$\mathrm{P}[\exists P_{st} \subseteq G^{CF}(D_n), |P_{st}| = r \mid d_s, d_t] = \frac{d_s d_t}{m} O(E[D_n^2]^{r-1}).$$

where m is the number of unmatched stubs.

Proof. Let $G = G^{\mathrm{CF}}(D_n)$. By $M(n) := (n-1)!!$ we denote the number of matchings on n vertices, where !! denotes the double factorial:

$$n!! := \begin{cases} n \cdot (n-2) \cdot \ldots \cdot 5 \cdot 3 \cdot 1 & \text{for } n > 0 \text{ odd,} \\ n \cdot (n-2) \cdot \ldots \cdot 6 \cdot 4 \cdot 2 & \text{for } n > 0 \text{ even, and} \\ 1 & n \in \{0, -1\}. \end{cases}$$

We will need the following bound for $k < n$:

$$\frac{M(n-k)}{M(n)} \leqslant \Big(\frac{(2e)^k (n-k)^{n-k}}{n^n} \Big)^{1/2} \leqslant \Big(\frac{2e}{n} \Big)^{k/2}.$$

The number of available stubs decreases with each edge added to the graph and hence the probability of an edge crucially depends on the number m of *remaining* stubs.

Fix a path P_{st} of length r and let $d_1, \ldots, d_{r-1}$ denote the weights of its internal vertices. The probability of this path existing in G, conditioned on the weight of its endpoints, is bounded by

$$\begin{aligned} \mathrm{P}[P_{st} \subseteq G \mid d_s, d_t] &\leqslant d_s d_t \frac{M(m-2r)}{M(m)} \sum_{d_1, \ldots, d_{r-1}} \prod_{i=1}^{r-1} d_i^2 \, \mathrm{P}[\hat{D}_n = d_i] \\ &\leqslant \frac{d_s d_t}{m^r} O(E[D_n^2]^{r-1}). \end{aligned}$$

Therefore the probability that *some* s-t-path of length r exists is

$$\mathrm{P}[\exists P_{st} \subseteq G \mid d_s, d_t] \leqslant \frac{d_s d_t}{m} O(E[D_n^2]^{r-1}),$$

as claimed. □

Proof of Theorem 3 for $G^{CF}(D_n)$. By Lemma 13, the probability of an s-t-path of length r existing in $G := G^{\mathrm{CF}}(D_n)$ is

$$\mathrm{P}[\exists P_{st} \subseteq G^{\mathrm{CF}}(D_n), |P_{st}| = r \mid d_s, d_t] = \frac{d_s d_t}{m} O(E[D_n^2]^{r-1}).$$

Note that this probability looks almost identical to the one given by Lemma 12, provided that the number of remaining stubs m is $\Theta(n)$. Since we want to estimate the probability of the event $\mathcal{E}^X_{r,\xi}$, only up to $r\xi|X|$ edges need to be considered at once; meaning that at least

$$\hat{m} - 2r\xi|X| \geqslant 2\mu n - 2r\xi n/4e(r\xi+1) \geqslant (\mu-1)2n$$

stubs remain (again we let $\mu = E[D_n]$). By coupling we see that

$$\mathrm{P}[\mathcal{E}^X_{r,\xi}]_{G^{\mathrm{CF}}(D_n)} \leqslant \mathrm{P}[\mathcal{E}^X_{r,\xi}]_{G^{\mathrm{CL}}(\Theta(D_n))}$$

and we conclude that $G^{\mathrm{CF}}(D_n)$ has bounded expansion for distributions with tail-bound $\Omega(d^{3+\varepsilon})$. Similarly, the event $\mathcal{K}^X_r$ for any set of vertices $|X| \leqslant \sqrt{n/4er}$ concerns at most $n/4e$ edges and hence the number of stubs left is $m = \Theta(n)$. Thus

$$\mathrm{P}[\mathcal{K}^X_r]_{G^{\mathrm{CF}}(D_n)} \leqslant \mathrm{P}[K^X_r]_{G^{\mathrm{CL}}(\Theta(D_n))}$$

and therefore $G^{\mathrm{CF}}(D_n)$ is nowhere dense for distributions with a tail that is in $\Theta(d^{3+o(1)})$. The lower bounds provided by Lemma 10 and Lemma 11 can be easily adapted in a similar way to apply to the configuration model. □

3.6 Perturbed bounded-degree graphs

An interesting application of the well-understood Erdős-Rényi random graphs is as a *perturbation* of some n-vertex base graph $G^\star$. We will use the notation $G = G^\star \cup G(n, \mu/n)$ to denote the graph obtained from $G^\star$ by adding in every possible edge not already contained in $G^\star$ independently with probability μ/n. Obviously, the graph $G^\star$ itself might be a random graph itself.

This procedure is more flexible than many existing generalizations of the Erdős–Rényi model, like (sparse) *inhomogeneous random graphs* [13] or (sparse) *generalized random graphs* [17, 78]. Similar (e.g. [54]) and more general (e.g. [7]) models have been defined before, yet there seems to be no consensus on a name or notation.

In particular, uniform perturbation can be seen as the baseline for more complicated models, like the small-world model by Kleinberg (described below), models used in percolation theory (e.g. [12]) and the hybrid model by Chung and Lu [23] (described above). The central question is: what graph classes are still structurally sparse after the addition of few random edges?

We will call the graph G drawn in the first step the *base graph*. In the Erdős–Rényi model, the base graph G would be the edgeless graph and the edge probabilities p_n constant functions for all n.

Theorem 4. *Let $\mathcal{G}$ be a class of bounded-degree graphs and μ a constant. Then the composite model $G^U(\mathcal{G}) + G(n, \mu/n)$ has bounded expansion with high probability.*

Note that this theorem in particular applies to $G(n, \mu/n)$ itself. The result carries over to the *stochastic block model*, if the parameters involved are small enough. This model was first studied in mathematical sociology by Holland, Laskey, and Leinhardt in 1983 [42] and extended by Wang and Wong to directed graphs [96]. We will supplement the above result by demonstrating that there exist very sparse graph classes of unbounded degree for which such a perturbation results in dense clique minors.

The following technical lemma subsumes Theorem 4.

Lemma 14. *Let $\mathcal{G}$ be a class of graphs with the following properties:*

- *$\mathcal{G}$ has bounded expansion, and*
- *for $G \in \mathcal{G}$ and every $r \in \mathbf{N}$ the distribution of $|N^r|$ has a tail-bound h with $h(d) = \Omega(d^{3+\varepsilon})$ for some $\varepsilon > 0$.*

Then $\mathcal{G} + G(n, \mu/n)$ has bounded expansion with high probability.

Proof. Let $G_\nabla \in \mathcal{G}$, $\widetilde{G} = \mathcal{G}(n, \mu/n)$ and let $G = G_\nabla + \widetilde{G}$. Assume H is an r-shallow topological minor of G and consider an embedding ϕ_V, ϕ_E of H witnessing this fact. Since $\widetilde{\nabla}_r(G_\nabla)$ is a constant, most of H's density must depend on random edges, i.e.

$$|\{e \in H \mid \phi_E(e) \cap E(\widetilde{G}) = \varnothing\}| \leqslant \widetilde{\nabla}_r(G_\nabla)|H|.$$

Therefore it suffices to bound the density of topological minors whose embedding use at least one edge of $\widetilde{G}$ for each edge of the minor. Consider a path P of length r in G that uses at least one edge of $\widetilde{G}$: each component of $P \setminus E(\widetilde{G})$ is contained in a subgraph $G_\nabla[N^r(v)]$ for some vertex v. Let $N_1, N_2, \ldots, N_p$ be these subgraphs of the path P: then we can bound the probability that P exists by considering the probability that there exist at least one edge between N_i and N_{i+1} in $\widetilde{G}$, for $1 \leqslant i \leqslant p-1$.

Define the random variable $D_{r,G}$ with distribution given by

$$\Pr[D_{r,G} = d] = \frac{|\{v \in G : |N^r(v)| = d\}|}{|G|}.$$

Since the probability that two r-neighbourhoods $N^r(u)$, $N^r(v)$ in G_∇ are connected by an edge in $\widetilde{G}$ is at most

$$\frac{\mu |N^r(u)||N^r(v)|}{n}$$

we can couple the occurrence of r-paths in G to the occurrence of edges in $G^{\mathrm{CL}}(D_{r,G})$. Hence we have that

$$\widetilde{\nabla}_r(G) - \widetilde{\nabla}_r G_\nabla \leqslant \widetilde{\nabla}_r(G^{\mathrm{CL}}(D_{r,G}))$$

in the stochastic sense. Since the latter has bounded expansion with high probability by Theorem 3, we conclude that $\mathcal{G} + G(n, \mu/n)$ does as well. □

The above result poses the question: are there structurally sparse classes which do *not* stay sparse under perturbation? The answer is yes: consider the class of graphs consisting of $\Theta(\sqrt{n})$ copies of $S_{\sqrt{n}}$. The probability that two such stars will be connected by a randomly added edge is lower-bounded by some constant, hence the minor obtain by contracting the former stars has density $\Theta(n)$ (while only having $\sqrt{n}$ vertices). Hence, the perturbed class is actually somewhere dense with high probability.

This example can be easily generalised: the presence of n^α vertices X to which we can assign at least n^β of their respective r-neighbours (not assigning any neighbour to more than on vertex of X), for some constant r, will yield an r-shallow minor whose density is concentrated around $n^{3\alpha+\beta-1}$. Hence for all α, β that satisfy $3\alpha + \beta > 2$, the perturbed class is somewhere dense with high probability.

4 Graph Models without Bounded Expansion

In this section we consider the Kleinberg [48, 49] and Barabási-Albert [6, 10] Models, which, respectively, were designed to replicate "small-world" properties and heavy-tailed (power-law) degree distributions observed in complex networks. We show that both these models (with typical parameters) do not have bounded expansion, and in fact are somewhere dense w.h.p./non-vanishing probability, respectively. This is done by showing the existence of two/one-subdivisions of cliques respectively in the generated graphs with a certain probability.

4.1 The Kleinberg Model

Many social networks exhibit a property that is commonly referred to as the "small-world phenomenon." This property asserts that any two people in a network are likely to be connected by a short chain of acquaintances. This was first observed by Stanley Milgram in a study published in 1967 [67]. Milgram's study suggested that individuals in a social network who only knew the locations of their immediate acquaintances are collectively able to construct short chains between two points in the network. More recently,

Kleinberg proposed a family of network models to explain the success of decentralized algorithms in finding short paths in social networks [49].

Kleinberg's model starts with a $n \times n$ grid as the base graph and allows edges to be directed. For a universal constant $p \geqslant 1$, a node u has a directed edge to every other node within lattice distance p. These are the *local neighbors* of u. For universal constants $q \geqslant 0$ and $r \geqslant 0$, node u has q *long range neighbors* chosen independently at random. The i^{th} directed outarc from u has endpoint v with probability $d(u,v)^{-r} / \sum_x d(u,x)^{-r}$.

When $r = 0$, the long-range contacts are uniformly distributed throughout the grid, and one can show that there exist paths between every pair of nodes of length bounded by a polynomial in $\log n$, exponentially smaller than the number of nodes. Kleinberg shows that in this case, the expected delivery time of *every* decentralized algorithm (one that uses only local information) is $\Omega(n^{2/3})$. When $p = q = 1$ and $r = 2$, then short chains continue to exist between the nodes of the network, but here is a decentralized algorithm to transmit a message that takes $O(\log^2 n)$ time in expectation between any two randomly chosen points.

What Kleinberg's model shows is that if the long-range contacts are formed independently of the geometry of the grid, then short chains exist between every pair of nodes, but nodes working with local knowledge are unable to find them. If the long-range contacts are formed by taking into account the grid structure in a specific way, then short chains exist and nodes working with local knowledge are able to discover them. We show that for those parameters where greedy routing is efficient, not only does the model not have bounded expansion, it is in fact, *somewhere dense* w.h.p.

Theorem 5. *The Kleinberg model with parameters $p = q = 1$ and $r = 2$ is somewhere dense w.h.p.*

Proof. Let Γ_n be an $n \times n$ grid. For $p = q = 1$ and $r = 2$, the probability that a node u has v as its long-range contact is proportional to $d_{\Gamma_n}(u,v)^{-1}$ and the normalizing factor in this case is $O(1/\log n)$. This can be easily seen by summing up $1/d_{\Gamma_n}(u,x)^2$ for all x and noticing that in the grid, there are $4d$ neighbors that are at a distance of d from u.

$$\sum_x \frac{1}{d_{\Gamma_n}(u,x)^2} = \sum_{d=1}^{n} \frac{4d}{d^2} \approx 4 \log n.$$

To show that the model is somewhere dense, we show that 2-subvidisions of cliques of a certain size $g(n)$ occur with high probability. Later we will see that $g(n) = \Omega(\log \log n)$. To this end, let $\Gamma'_{c \cdot g(n)}$ denote some fixed $c \cdot g(n) \times c \cdot g(n)$ subgrid of Γ_n, where c is some constant that we will fix later. Choose V' and

E' to be, respectively, a set of $g(n)$ nodes and a set of $g(n)^2$ edges from the subgrid $\Gamma'_{c\cdot g(n)}$ with the following properties: (i) the endpoints of the edges in E' are different from the nodes in V'; (ii) no two edges in E' share an endpoint.

Given any pair of vertices $u, v \in V'$ and an edge $e \in E'$ with endpoints a, b, the probability that a has u as its long-range neighbor is $\Omega((d_{\Gamma_n}(a,u) \cdot \log n)^{-1})$. Similarly, the probability that b has v as its long-range neighbor is $\Omega((d_{\Gamma_n}(b,v) \cdot \log n)^{-1})$ The probability of both these events happening is

$$\frac{1}{d_{\Gamma_n}(a,u)^2 d_{\Gamma_n}(b,v)^2} \cdot \frac{1}{\log^2 n} \geqslant \frac{1}{c^4 g(n)^4 \log^2 n}, \tag{1}$$

where we upper-bounded distances $d_{\Gamma_n}(x,y)$ by $c \cdot g(n)$. Thus the probability that there exists a 2-subdivided $g(n)$-clique in $\Gamma'_{c\cdot g(n)}$ is at least:

$$\left(\frac{1}{c^4 g(n)^4 \log^2 n}\right)^{g(n)^2} =: f(n,c). \tag{2}$$

The probability that there does *not* exist a 2-subdivided $g(n)$-clique in $\Gamma'_{c\cdot g(n)}$ is at most $1 - f(n,c)$. Hence the probability that there does not exist a 2-subdivided $g(n)$-clique in *any* $c \cdot g(n) \times c \cdot g(n)$ subgrid is at most

$$(1 - f(n,c))^{\frac{n}{c^2 g(n)^2}} \leqslant \exp\left(-\frac{n}{c^2 \cdot g(n)^2 \cdot \left(c^4 \cdot g(n)^4 \cdot \log^2 n\right)^{g(n)^2}}\right) := e^{\frac{n}{h(n)}}.$$

This follows from the inequality $(1 - x/p)^p \leqslant e^{-x}$.

Choose $g(n) = \log\log n$ and $c = 3$ (actually any $c \geqslant 3$ works.) Then it is easy to show that $h(n) < \sqrt{n}$. Thus the probability of a 2-subdivided $(\log\log n)$-clique *not* existing is at most $e^{-\sqrt{n}}$ (which goes to zero as $n \to \infty$), and we conclude that the graph model is somewhere dense. □

4.2 The Barabási-Albert model

The Barabási-Albert model uses a preferential attachment paradigm to produce graphs with a degree distribution that mimics the heavy-tailed distribution observed in many real-world networks [10]. This model uses a random graph process that works as follows: Start with a small number n_0 of nodes and at every time step, add a new node and link it to $n \leqslant n_0$ nodes already present in the "system." To model preferential attachment, we assume that the probability with which a new node u is connected to node v already present in the system is proportional to the degree of v, so that $\mathrm{P}[u \to v] = \deg(v) / \sum_x \deg(x)$, where the sum in the denominator is over all

vertices x that are already in the system. After t time steps, the model leads to a random network with $t + n_0$ vertices and nt edges.

Barabási and Albert suggested that such a network evolves into one in which the fraction $P(d)$ of nodes of degree d is proportional to $d^{-\gamma}$. They observed experimentally that $\gamma = 2.9 \pm 0.1$ and suggested that γ is actually 3. This model was rigorously analyzed by Bollobás, Riordan, Spencer and Tusnády in [15] who showed that it is indeed the case that the fraction of vertices of degree d fall off as d^{-3} as $d \to \infty$. In [14], Bollobás and Riordan showed that the diameter of the graphs generated by this model is asymptotically $\log n / \log \log n$.

We first provide a formal restatement of the Barabási-Albert Model. Note that this is slightly different from the formalization of Bollobás et al. in [15]. We start with a "seed" graph G_0 with n_0 nodes $u_1, \ldots, u_{n_0}$ with degrees $d_1, \ldots, d_{n_0}$. The number of edges in the seed graph is denoted by m_0. At each time step $t = 1, 2, \ldots$, we create a graph G_t by adding a new node v_t and linking it to q nodes in G_{t-1}. These q nodes are picked independently and with a probability that is proportional to their degrees in G_{t-1}. That is, we choose $u \in V(G_{t-1})$ to link to with probability

$$\mathrm{P}\big[\{v_t, u\} \in E(G_t)\,\big] = \frac{\deg_{G_{t-1}}(u)}{2 \cdot |E(G_{t-1})|} = \frac{\deg_{G_{t-1}}(u)}{2(m_0 + q(t-1))}.$$

Note that in this model, all edges between v_t and nodes of G_{t-1} are assumed to be added simultaneously (so that the increasing degrees of nodes which receive edges from v_t do not influence the probabilities for this time step.) In this restated version, n_0, $d_1, \ldots, d_{n_0}$, and q are the parameters of the model.

Lemma 15. *Given any fixed r, a graph G_n generated by the preferential attachment model with parameters n_0, $d_1, \ldots, d_{n_0}$ and $q \geqslant 2$ has a 1-subdivided K_r as a subgraph with probability at least $(4(\frac{m_0}{q} + r + r^2))^{-r^2}$, provided $n \geqslant r + r^2$.*

Proof. Choose any $r \in \mathbf{N}$. We will show that there exists a 1-subdivided K_r in the graph with probability that depends only on m_0, q, and r.

Consider the graph after the first $r + r^2$ time steps. Let $v_1, \ldots, v_r$ be the nodes that were added in the first r time steps and fix two nodes v_i, v_j from among these. The probability that a new node v_k (for any $r + 1 \leqslant k \leqslant r^2$) is connected to these two fixed nodes is at least

$$\left(\frac{q}{2(m_0 + q(r + r^2))}\right)^2 =: f(m_0, r, q),$$

where the denominator is the sum of the vertex degrees after $r + r^2$ time steps. Now if v_{r+1} is linked to v_1, v_2 and v_{r+2} is linked to v_1, v_3 and so on such

that the nodes added after time step r connect the first r nodes in a pairwise fashion, we would have a 1-subdivided K_r in the graph G_{r+r^2}. The probability of this happening is at least $f(m_0, r, q)^{r^2}$. Thus, for every r, the probability of 1-subdivided K_r existing is non-zero if the graph is large enough. □

It immediately follows that the Barabási-Albert model (and similar preferential attachment models) is not a.a.s. nowhere dense and in particular does not have bounded expansion a.a.s. We note that this result is more of theoretical interest, since the probabilities involved might be small enough to be irrelevant in practice. As such it would be worthwhile to investigate whether the Barabási-Albert model is somewhere dense a.a.s.

5 Experimental Evaluation

Although it has been established that real-world networks are sparse, and tend to have low degeneracy (relative to the size of the network), it is natural to ask whether there is empirical evidence that they satisfy the stronger conditions of bounded expansion. Unfortunately, since bounded expansion itself is a property of graph classes and not single graphs, it is impossible to determine whether individual instances are bounded expansion or not. One natural proxy would be to evaluate the *grad* of these graphs, calculating the maximum density of an r-shallow minor for each $r \in \mathbf{N}$ (obviously stopping when r is the diameter of the network), but it is not known how to find such minors in reasonable time.

In order to get around these difficulties, we calculate upper bounds on χ_{p-1} (the p-centered coloring number), a good proxy, since it and the grad are both related to each other by factors independent of the graph size. This is further justified by the fact that p-centered colorings are directly applicable to algorithm design, where the complexity of such algorithms depends heavily on the number of required colors (as will be shown in Section 6). Since it is very time-consuming to obtain a good p-centered coloring for large p (this is analogous to determining a reasonable bound for the maximum density of an r-shallow minor for large r), we evaluate this property for small values only. This is also roughly the range of p which is relevant to the algorithms presented later in this paper when applied to practical settings.

To obtain upper bounds on χ_{p-1}, we implemented the transitive- fraternal augmentation procedure of Nešetřil and Ossona de Mendez [72]. Our theoretical results predict that graphs generated with the configuration model for typical degree distributions of complex networks will likely have bounded expansion. We thus compared the results for χ_3 of this procedure between real networks and networks generated using the configuration model for the

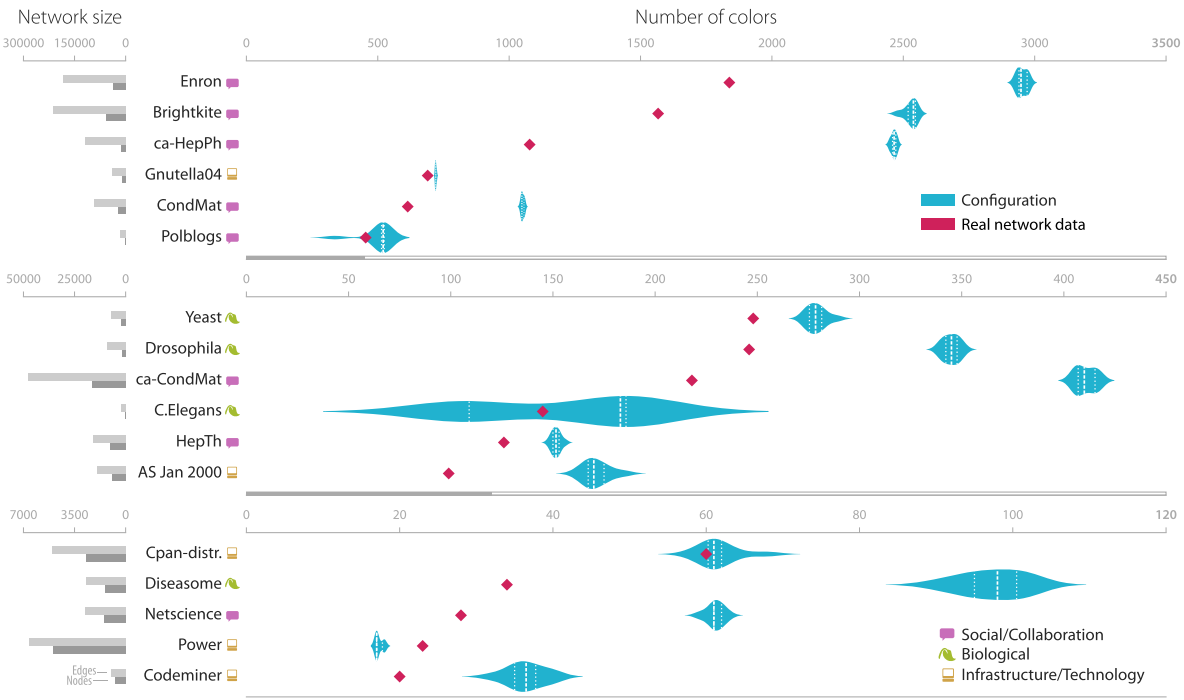

Figure 1: Comparison of 4-centered coloring numbers on real-world networks (red diamonds) compared to synthetic graphs (blue violins) with the same degree distribution. Each violin represents 10 random instances generated with the configuration model, with median and quartiles marked with dashed and dotted lines. Networks are partitioned into three groups by size (indicated on the left) to enable rescaling axes. See Table 2 for data sources.

same degree distributions. We chose χ_3 since it is relatively easy to compute for large networks but still heavily influenced by one-subdivisions of cliques. The results of this experiment can be found in Figure 1.

For almost all networks the bound for the real-world network is either smaller or comparable to the values for the synthetic graphs. The one exception is "power", which is not surprising since this network has a relatively complex (grid-like) structure over low degree nodes.

We furthermore extended the algorithm based on transitive-fraternal augmentations with simple heuristic improvements (e.g. giving high degree nodes their own private color, merging color classes where possible), ran it to find p-centered colorings on a small corpus of well-known complex networks and verified the results. The results of the best colorings we were able to achieve with this relatively simple methos can be found in Table 2.

The results show that some networks clearly have a moderately growing grad; in particular the larger networks *Netscience*, both *Cpan-* networks and *Diseasome*. Other networks, like *Twittercrawl*, have such quickly growing

			p					
Network	Vertices	Edges	2	3	4	5	6	∞
Karate [98]	35	78	6	7	9	9	10	8
Dolphins [61]	62	159	7	11	17	18	19	24
Lesmiserables [51]	77	254	10	15	16	16	16	16
Polbooks [53, 76]	105	441	8	16	22	29	31	30
Word-adjacencies [80]	112	425	8	18	27	35	41	48
Football [38]	115	613	9	22	33	49	62	69
Airlines [2]	235	1297	11	28	39	47	55	64
Sp-data-school [92]	238	5539	23	100	138	157	168	171
C.Elegans [97]	306	2148	8	36	74	83	118	153
Hex-grid	331	930	3	9	20	21	25	69
Codeminer [2]	724	1017	5	10	15	17	23	51
Cpan-authors [1]	839	2212	9	24	34	43	47	224
Diseasome [39]	1419	2738	12	17	22	25	30	30
Polblogs [3]	1491	16715	30	118	286	354	392	603
Netscience [80]	1589	2742	20	20	28	28	28	20
Drosophila [93]	1781	8911	12	65	137	188	263	395
Yeast [18]	2284	6646	12	38	178	254	431	408
Cpan-distr. [1]	2719	5016	5	14	32	42	56	224
Twittercrawl [2]	3656	154824	89	561	1206	1285	1341	–
Power [97]	4941	6594	6	12	20	21	34	95
AS Jan 2000 [56]	6474	13895	12	29	70	102	151	357
Hep-th [77]	7610	15751	24	25	104	328	360	558
Gnutella04 [56, 89]	10876	39994	8	43	626	–	–	–
ca-HepPh [56]	12008	118489	239	296	1002	–	–	–
CondMat [77]	16264	47594	18	47	255	1839	–	1310
ca-CondMat [56]	23133	93497	26	89	665	–	–	–
Enron [50, 59]	36692	183831	27	214	1428	–	–	–
Brightkite [20]	58228	214078	39	193	1421	–	–	–

Table 2: Number of colors in p-centered colorings computed on real world networks and upper bounds of their respective treedepth. These networks were mostly taken from the datasets found in [2, 57, 76].

p-centered coloring numbers that we did not invest the time to determine the value for larger p. Since graphs of bounded crossing number (as infrastructure networks tend to be) and bounded degree have bounded expansion, we are not surprised at the small number of colors needed by *Power* and *Hex*. Finally some networks, like *CondMat* and *Hep-th*, start out reasonably well for small p but show a sudden jump at $p = 3$. At present we do not know whether this is an artifact of the procedure we use to obtain the coloring or whether the

networks have indeed dense minors from a certain depth on. As shown in Section 4, some complex network models predict such an occurrence already for depth at most two. The growth behavior for small p as depicted in Table 2 might therefore serve as a property to distinguish types of networks, meriting future research.

In the last column, we provide upper bounds for the treedepth of these graphs. Notice that the number of colors needed for a p-centered coloring will be fewer than the treedepth of the the graph for any $p \leqslant n$: Given a treedepth decomposition of a graph of depth t we can color every node by its depth in the treedepth decomposition. Since the treedepth is a hereditary property, the graph induced by any subset of nodes will have treedepth at most t. Notice that the simple coloring algorithm we used sometimes colors the graph with more than t colors. This is a good indication that a better coloring algorithm or heuristic exists.

We also note that the subgraph isomorphism algorithm presented in Section 6 should be directly applicable to some of these graphs, given their comparatively low treedepth.

Finally, we argue that for practical purposes the definition of low treedepth colorings is too strict. Remember that per definition, in a p-centered coloring every graphs induced on $i < p$ colors has treedepth $\leqslant i$. Relaxing this latter condition, we arrive at the following variation of Proposition 1:

Observation 2. *Let $\mathcal{G}$ be a graph class of bounded expansion. There exists function f and g such that for every $G \in \mathcal{G}$, $p \in \mathbf{N}$, the graph G can be colored with $f(r)$ colors so that any $i < p$ color classes induce a graph of treedepth $\leqslant g(i)$ in G.*

Obviously this follows from Proposition 1, taking g as the identity. However, algorithms based on low treedepth colorings would run faster if g allows for a larger margin, provided we can decrease the number of colors f. This is owed to the fact that for large number of colors, iterating through all $< p$-sized subsets will be the deciding factor in the running time.

6 Algorithms

In this section, we present efficient algorithms for several important problems arising in the study complex networks which exploit bounded expansion. Note that in both cases, only the algorithm's running time relies on the grad being small, not its correctness—in that sense, these algorithms are oblivious to whether the input graph is sparse or not (unless, for example, algorithms that rely on planarity) The problems that we discuss revolve around the themes of *subgraph counting* and *centrality estimation*.

Computing the frequency of small fixed pattern graphs inside a network is the key algorithmic challenge in using *network motifs* and *graphlet degree distributions* to analyze network data (both of which are described in more detail in Section 6.1). We present a parameterized algorithm for counting the number of subgraphs with at most h vertices with a running time of $6^h \cdot t^h \cdot n$, where t is the treedepth of the input graph. In a graph class of bounded expansion, we use this algorithm in conjunction with p-centered colorings.

Another topic of interest in complex networks is estimating the relative importance of a vertex in the network (for example, how influential a person is inside a social network, which roads are busiest in a road-network, or which location is most attractive for business). The typical approach is to define/select an appropriate *centrality measure* (see [52] for a survey of common measures). We focus on the *closeness centrality*, which was introduced by Sabidussi [90], and related extensions. These measures are related in the sense that computing them requires knowledge of all the pairwise distances between the vertices of the network, which even in sparse networks takes time $O(n^2)$ to compute [16]. We introduce localized variants of the closeness-based centrality measures and design a linear-time algorithm to compute them in bounded expansion classes and provide experimental data that suggests that these measures are able to recover the topmost central vertices quite well.

We will make extensive use of the following proposition to compute the localized centrality measures.

Proposition 4 (Truncated distances [72])**.** *Let G be a graph of bounded expansion. For every r one can compute in linear time a directed graph $\vec{G}_r$ with in-degree bounded by $f(r)$ – for some function f – on the same vertex set as G and an arc labeling $\omega\colon \vec{E}(\vec{G}_r) \to \mathbf{N}$ such that for every pair $u, v \in G$ with $d_G(u,v) \leqslant r$ one of the following holds:*

(i) $uv \in \vec{G}_r$ *and* $\omega(uv) = d_G(u,v)$*;*

(ii) $vu \in \vec{G}_r$ *and* $\omega(vu) = d_G(u,v)$*;*

(iii) there exists $w \in N^-_{\vec{G}_r}(u) \cap N^-_{\vec{G}_r}(v)$ *such that* $\omega(wu) + \omega(wv) = d_G(u,v)$*.*

6.1 Counting graphlets and subgraphs

In the following we highlight three domain-specific applications of computing the frequency of small fixed pattern graphs inside a network. In particular, the concept of *network motifs* and *graphlets* has proven very useful in the area of computational biology.

A *network motif* is a subgraph (not necessarily induced and possibly labeled) that appears with a significantly higher frequency in a real-world

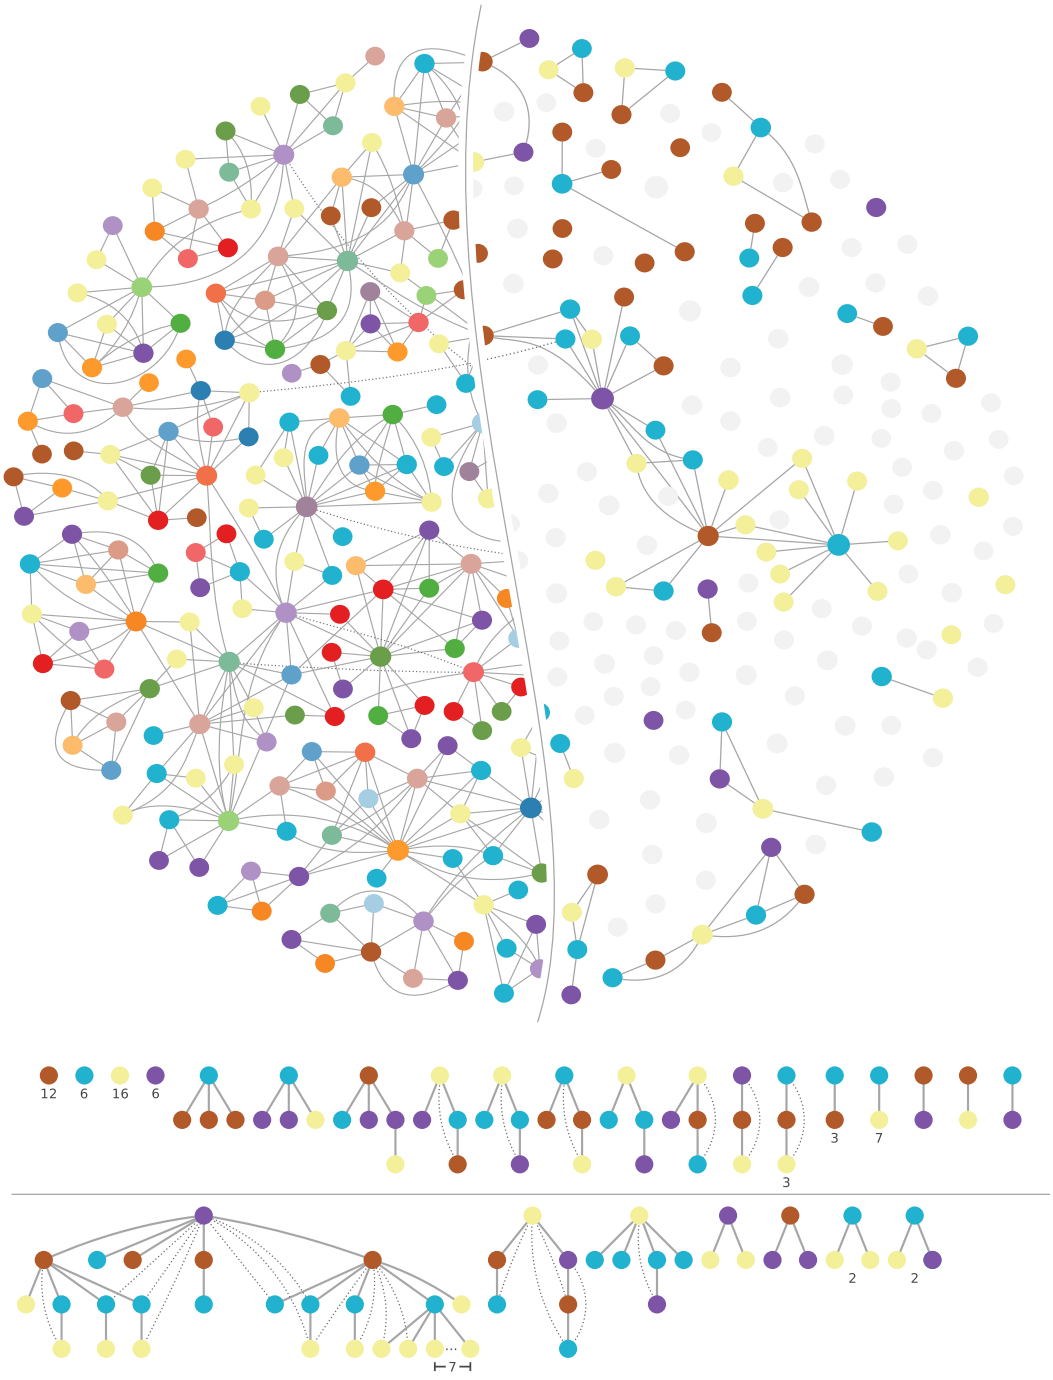

Figure 2: A 5-centered coloring (using 21 colors) of a real-world social network, Newman's Network Science [80] (giant component shown), which represents co-authorships between researchers in the field of network science as of 2006. The right half is restricted to a a subgraph formed by 4 color classes. Below, the corresponding representation by trees of depth $\leqslant 4$ (with multiplicities noted).

network than one would expect by pure chance. Introduced in [68] under the hypothesis that such frequently occurring structures have a functional significance, motifs have been identified in a plethora of different domains—including protein-protein-interaction networks [5], brain networks [91] and electronic circuits [43]. We point the interested reader to the surveys of Kaiser, Ribeiro and Silva [88] and Masoudi-Nejad, Schreiber, and Kashani [63] for a more extensive overview.

Graphlets are a related concept, though their application is in an entirely different scope. While motifs are used to identify and explain local structure in networks, graphlets are used to 'fingerprint' them. Pržulj [86] introduced the *graphlet degree distribution* as a way of measuring network similarity. To compute it, one enumerates all connected graphs up to a fixed size (five in the original paper) and computes for each vertex of the target graph how often it appears in a subgraph isomorphic to one of those patterns. Since some graphlets exhibit higher symmetry than others, the computation takes into account all possible automorphisms. The degree distribution then describes for each graphlet G_i, how many vertices of the target graph are contained in $0, 1, 2, \ldots$ subgraphs isomorphic to G_i—more precisely, in how many orbits of the respective automorphism groups it appears in. Note that if the set of graphlets only contains the single-edge graph this computation yields exactly the classical degree distribution.

The application of this distribution is two-fold: On the one hand, it can be used to measure similarity of multiple networks, in particular, networks related to biological data [41]. On the other hand, the local structure around a vertex can reveal domain-specific functions. This is the case for protein-protein interaction networks, where local structure correlates with biological activity [66]. This fact has been applied to identify cancer genes [65] and construct phylogenetic trees [55]. Graphlets have further been used in analysis of workplace dynamics [94], photo cropping [19] and DoS attack detection [85].

A third application of subgraph counting was given by Ugander et al. [95]: their empirical analysis and subsequent modeling of social networks revealed that there is an inherent bias towards the occurrence of certain subgraphs. Thus frequencies of small subgraphs are an important indicator for the social domain, similar to the role of graphlet frequencies in biological networks.

In Theorem 18.9 from [73] it was shown that for a graph class of bounded expansion counting the number of satisfying assignments of a fixed boolean query is possible in linear time on a labeled graph. This implies that (labeled) graphlet and motif counting are linear time on a graph classes of bounded expansion. This result is achieved by using the algorithm presented in Lemma 17.3 in [73] to count the number of satisfying assignments of a fixed boolean query parameterized by treedepth. The running time of this algorithm

is $O(2^{ht} \cdot ht \cdot n)$, where h is the number of nodes in the graphlet or motif and t is the depth of a treedepth decomposition of the graph. We provide an algorithm with a running time of $O(6^h \cdot t^h \cdot h^2 \cdot n)$. This achieves a better running time when used to count on graphs of bounded expansion, since then, as will be explained later, h will equal t, i.e. a constant. We discuss how with a small modification, this algorithm can also count how many times a node appears as a specific node of a specific graphlet or motif.

The tool of choice for applying a counting algorithm designed for bounded-treedepth graphs to a class of bounded expansion is a p-centered coloring: to compute the frequency of a given pattern of size k, we compute a $(k+1)$-centered coloring of the input graph in linear time as per Proposition 1. We then enumerate all possible choice of $i < k$ colors and count the frequency of the pattern graph in the graph induced by those colors classes. As this induced subgraph has bounded treedepth, we can focus on counting a fixed subgraph inside a target graph of treedepth at most k. It is then easy to compute the frequency for the original graph using inclusion-exclusion on the color classes.

Central to the dynamic programming we will use to count isomorphisms is the following notion of a *k-pattern* which is very similar to the well-known notion of *boundaried graphs*.

Definition 12. *A* k-pattern *of the graph H is a triple $M = (W, X, \pi)$ where $X \subseteq W \subseteq V(H)$, $|X| \leqslant k$, such that $W \setminus X$ has no edge into $V(H) \setminus W$, and $\pi \colon X \to [k]$ is an injective function. We will call the set X the* boundary *of M. For a given k-pattern M we denote the underlying graph by $H[M] = H[W]$, the vertex set by $V(M) = W$, the boundary by $\mathrm{bd}(M) = X$ and the mapping by π^M.*

We denote by $\mathcal{P}_k(H)$ the set of all k-patterns of H. Note that every k-pattern (W, X, π) is also a $(k+1)$-pattern. In the following we denote by $|H| = |V(H)|$.

Lemma 16. *Let H be a graph. Then $|\mathcal{P}_k(H)| \leqslant 3^{|H|} \cdot k^{|H|}$.*

Proof. The vertices of H can be partitioned in $3^{|H|}$ possible ways into boundary vertices, pattern vertices and remainder. The number of ways an injective mapping for a boundary of size $b \leqslant |H|$ into $[k]$ can be chosen is bounded by $k^{|H|} = 2^{|H| \log k}$. In total the size of $\mathcal{P}_k(H)$ is always less than $3^{|H|} \cdot k^{|H|}$. □

The following definition show how k-patterns will be used structural, namely by gluing them together or by demoting a boundary-vertex to a simple vertex. These operations will later be used for dynamic programming.

Definition 13 (k-pattern join)**.** *Let H be a graph and $M_1 = (W_1, X_1, \pi_1)$, $M_2 = (W_2, X_2, \pi_2)$ k-patterns of H. Then the two patterns are* compatible *if*

$W_1 \cap W_2 = X_1 = X_2$ *and for all* $v \in X_1$ *it holds that* $\pi_1(v) = \pi_2(v)$*. Their* join *is defined as the* k*-pattern* $M_1 \oplus M_2 = (W_1 \cup W_2, X_1, \pi_1)$*.*

Definition 14 (k-pattern forget)**.** *Let* H *be a graph, let* $M = (W, X, \pi)$ *be a* k*-pattern of* H *and* $i \in [k]$*. Then the* forget operation *is the* k*-pattern*

$$M \ominus i = \begin{cases} (W, X \setminus \pi^{-1}(i), \pi|_{X \setminus \pi^{-1}(i)}) & \text{if } \pi^{-1}(i) \neq \varnothing \text{ and } N_H(\pi^{-1}(i)) \subseteq W \\ \bot & \text{if } \pi^{-1}(i) \neq \varnothing \text{ and } N_H(\pi^{-1}(i)) \not\subseteq W \\ (W, X, \pi) & \text{otherwise} \end{cases}$$

Structurally, the k-pattern's boundaries will represent vertices from the path of the root vertex to the currently considered vertex in the treedepth decomposition, while the remaining vertices of the pattern represent vertices somewhere below it. The following two notations help expressing these properties.

Definition 15 (Subtree and root path)**.** *Let* T *be a treedepth decomposition of* G *rooted at* $r \in G$ *and let* $v \in V(G)$ *be a vertex. Then the* subtree of v *is the subtree* T_v *of* T *rooted at* v*. The* root path *of* v *is the unique path* P_v *from the root* r *to* v *in* T*. We let* $P_v[i]$ *denote the* i*th vertex of the path (starting at the root), so that* $P_v[1] = r$ *and* $P_v[\,|P_v|\,] = v$*.*

We can now state the main lemma. The proof contains the description of the dynamic programming which works bottom-up on the vertices of the given treedepth decomposition (i.e. starting at the leaves and working towards the root of the decomposition).

Lemma 17. *Let* H *be a fixed graph on* h *vertices. Given a graph* G *on* n *vertices and a treedepth decomposition* T *of height* t*, one can compute the number of isomorphisms from* H *to induced subgraphs of* G *in time* $O(6^h \cdot t^h \cdot h^2 \cdot n)$ *and space* $O(3^h \cdot t^h \cdot ht \cdot \log n)$*.*

Proof. We provide the following induction that easily lends itself to dynamic programming over T. Denote by $M_H = (V(H), \varnothing, \varepsilon)$ the trivial t-pattern of H, here $\varepsilon\colon \varnothing \to \varnothing$ denotes the null function. Consider a set of vertices $v_1, v_2, \ldots, v_\ell \in G$ with a common parent v in T with respective subtrees T_{v_i} and root paths P_{v_i} for $1 \leqslant i \leqslant \ell$. Note that the root paths $P_{v_1}, \ldots, P_{v_\ell}$ all have the same length k and share the path P_v as a common prefix.

Let M_1 be a fixed t-pattern of H. We define the mapping $\psi_v^{M_1}\colon \mathrm{bd}(M_1) \to P_v$ via $\psi_v^{M_1}(v) = P_v[\pi^{M_1}(v)]$ which takes the pattern's boundary and maps it to the vertices of the root-path.

For patterns M_1 that satisfy that for all $u \in \mathrm{bd}(M_1)$, $\pi^{M_1}[u] \leqslant l$, we denote by $f[v_1, \ldots, v_\ell][M_1]$ the number of isomorphisms $\phi_1\colon V(M_1) \to V(G)$ such that

(i) $\phi_1|_{\mathrm{bd}(M_1)} = \psi_v^{M_1}$

(ii) $\phi_1(V(M_1) \setminus \mathrm{bd}(M_1)) \subseteq G[V(T_{v_1} \cup \cdots \cup T_{v_\ell})]$.

In other words we charge subgraphs to patterns whose boundaries lie on the shared root-path P_v, such that the labeling of the boundary coincides with the numbering induced by P_v while the rest of the pattern is contained entirely in the subtree below v.

Let r be the root of the treedepth decomposition. By the above definition, $f[r][M_H]$ counts exactly the number of isomorphisms of H into subgraphs of G.

We will show now how we can compute $f[r][M_H]$ recursively. For a leaf $v \in T$ and a t-pattern $M_1 = (W_1, X_1, \pi_1) \in \mathcal{P}_k(H)$ we compute $f[v][M_1]$ as follows: Defined the value $p_v^{M_1}$ to be 1 if the function $\psi\colon W_1 \to P_v$ defined as $\psi(w) = P_v[\pi_1[w]]$ is an isomorphism from $H[W_1]$ to $G[\psi(W)]$ and 0 otherwise. In particular, $p_v^{M_1}$ will be zero if $W_1 \neq X_1$ or $|W_1| > P_v$. Then for the leaf v, we compute

$$f[v][M_1] = \sum_{M_2 \ominus |P_v| = M_1} p_v^{M_2}$$

where $M_2 \in \mathcal{P}_t(H)$.

The following recursive definitions show how $f[\cdot][M_1]$ can be computed for all inner vertices of T.

$$f[v][M_1] = \sum_{M_2 \ominus |P_v| = M_1} f[v_1, \ldots, v_\ell][M_2] \qquad \text{(forget)}$$

$$f[v_1, \ldots, v_{j-1}, v_j][M_1] = \sum_{M_2 \oplus M_3 = M_1} f[v_1, \ldots, v_{j-1}][M_2] \cdot f[v_j][M_3] \qquad \text{(join)}$$

where $M_2, M_3 \in \mathcal{P}_t(H)$.

We need to prove that the table f correctly reflects the number of isomorphisms to subgraphs satisfying properties i and ii.

Consider the *join*-case first: fix a pattern $M_1 \in \mathcal{P}_t(H)$. By induction, the entries $f[v_1, \ldots, v_{j-1}][\cdot]$ and $f[v_j][\cdot]$ correspond to the number of isomorphisms to subgraphs that satisfy properties i and ii with the node tuples $v_1, \ldots, v_{j-1}$ and v_j, respectively. We need to show that $f[v_1, \ldots, v_j][M_1]$ as computed above counts the number of isomorphisms from $H[M_1]$ to subgraph of G such that $\phi_1|_{\mathrm{bd}(M_1)} = \psi_v^{M_1}$ and $\phi_1(V(M_1) \setminus \mathrm{bd}(M_1)) \subseteq G[V(T_{v_1} \cup \cdots \cup T_{v_j})]$.

Consider the set Φ_1 of all isomorphisms from $H[M_1]$ to subgraphs of G satisfying properties i and ii for the vertex tuple $v_1, \ldots, v_j$. For any vertex subset $R \subseteq V(M_1) \setminus \mathrm{bd}(M_1)$, define the slice $\Phi_1(R) \subseteq \Phi_1$ as those isomorphisms ϕ that satisfy $\phi^{-1}(\phi(V(H)) \cap T_{v_j}) = R$. Let $L = (V(M_1) \setminus \mathrm{bd}(M_1)) \setminus R$ and define the patterns $M_L = (L \cup \mathrm{bd}(M_1), \mathrm{bd}(M_1), \pi^{M_1})$ and $M_R = (R \cup$

$\mathrm{bd}(M_1), \mathrm{bd}(M_1), \pi^{M_1})$. Then by induction $|\Phi_1(R)| = f[v_1, \ldots, v_{j-1}][M_L] \cdot f[v_j][M_R]$. Since $M_1 = M_L \oplus M_R$ and clearly $M_L, M_R \in \mathcal{P}_t(H)$, the sum computes exactly $\sum_{R \subseteq V(M_1) \setminus \mathrm{bd}(M_1)} |\phi_1(R)| = |\phi_1|$.

Next, consider the *forget*-case. Again, fix $M_1 \in \mathcal{P}_t(H)$ and let u be the parent of v in T. Let Φ_1 be the set of those isomorphisms from $H[M_1]$ to subgraphs of G for which $\phi_1|_{\mathrm{bd}(M_1)} = \psi_u^{M_1}$ and $\phi_1(V(M_1) \setminus \mathrm{bd}(M_1)) \subseteq G[V(T_v)]$. We partition Φ_1 into $\Phi_1 = \Phi_{1,v} \cup \Phi_{1,\bar{v}}$ where $\Phi_{1,v}$ contains those isomorphisms ϕ for which $\phi^{-1}(v) \neq \varnothing$ and $\Phi_{1,\bar{v}}$ the rest. Since $|\Phi_{1,\bar{v}}| = f[v_1, \ldots, v_\ell][M_1]$ we focus on $\Phi_{1,v}$ in the following. For $w \in V(M_1) \setminus \mathrm{bd}(M_1)$, define $\Phi_{1,v}(w)$ as the set of those isomorphisms ϕ for which $\phi(w) = v$. Clearly, $\{\Phi_{1,v}(w) \mid w \in V(M_1) \setminus \mathrm{bd}(M_1)\}$ is a partition of $\Phi_{1,v}$. Define the pattern $M_w = (V(M_1), \mathrm{bd}(M_1) \cup \{w\}, \pi_w^{M_1})$ where $\pi_w^{M_1}$ is π^{M_1} augmented with the value $\pi_w^{M_1}(v) = |P_v|$. Note that by construction $M_1 = M_w \ominus |P_v|$. By induction, $|\Phi_{1,v}(w)| = f[v_1, \ldots, v_\ell][M_w]$ and therefore

$$|\Phi_1| = |\Phi_{1,\bar{v}}| + \sum_{w \in V(M_1) \setminus \mathrm{bd}(M_1)} |\Phi_{1,v}(w)| = \sum_{M_2 \ominus |P_v|} f[v_1, \ldots, v_\ell][M_2]$$

It remains to prove the claimed running time. Initialization of f for a leaf takes time $O(|\mathcal{P}_t(H)|h^2)$ since we need to test whether the function ψ defined above is an isomorphism for each pattern $M_1 \in \mathcal{P}_t(H)$.

For the other vertices, a forget operation can be achieved in time $O(|\mathcal{P}_t(H)|)$ per vertex by enumerating all t-patterns, performing the forget operation and looking up the count of the resulting pattern in the previous table.

A join operation needs time $O(|\mathcal{P}_t(H)| \cdot h \cdot 2^h)$ per vertex, since for a given pattern M_1 those patterns M_2, M_3 with $M_1 = M_2 \oplus M_3$ are uniquely determined by partitions of the set $V(M_1) \setminus \mathrm{bd}(M_1)$.

In total the running time of the whole algorithm is $O(|\mathcal{P}_t(H)| \cdot 2^h \cdot h^2 \cdot n)$. Note that we only have to keep at most $O(t)$ tables in memory, each of which contains the occurrence of up to $|P_t(H)|$ patterns stored in numbers up to n^h. Thus in total the space complexity is $O(|\mathcal{P}_t(H)| \cdot t \cdot \log(n^h)) = O(|\mathcal{P}_t(H)| \cdot ht \cdot \log n)$. □

To count the occurrences of H as an induced subgraph instead the number of subgraph isomorphisms, one can simply determine the number of automorphisms of H in time $2^{O(\sqrt{h \log h})}$ [8, 64] and divide the total count by this value (since this preprocessing time is dominated by our running time we will not mention it in the following). Counting isomorphism to non-induced subgraphs can be done in the same time and space by changing the initialization on the leaves, such that it checks for an subgraph instead of an induced subgraph. Dividing again by the number of automorphisms gives the number of subgraphs. By allowing the mapping of the patterns to map several nodes

to the same value, we can use them to represent homomorphisms. Testing the leaves accordingly the same algorithm can be used to count the number of homomorphisms from H to subgraphs of G. By keeping all tables in memory, thus sacrificing the logarithmic space complexity, and using backtracking we can also label every node with the number of times it appears as a certain vertex of H.

From these observations and Lemma 16 we arrive at the following theorem:

Theorem 6. *Given a graph H on h vertices, a graph G on n vertices and a treedepth decomposition of G of height t, one can compute the number of isomorphisms from H to subgraphs of G, homomorphisms from H to subgraphs of G, or (induced) subgraphs of G isomorphic to H in time $O(6^h \cdot t^h \cdot h^2 \cdot n)$ and space $O(3^h \cdot t^h \cdot ht \cdot \log n)$.*

Note that for graphs of unbounded treedepth the running time of the algorithm degenerates to $O(6^h \cdot h^2 \cdot n^{h+1})$, which is comparable to the running time of $2^{O(\sqrt{h \log h})} \cdot n^h$ of the trivial counting algorithm.

Theorem 7. *Given a graph H and a a graph G belonging to a class of bounded expansion, there exists an algorithm to count the appearances of H as a subgraph of G in time*

$$O\left(\binom{f(h)}{h} \cdot 6^h \cdot h^{h+2} \cdot n\right)$$

where f is a function depending only on the graph class.

This immediately extends to nowhere dense classes, which have low treedepth-colorings with at most n^ε colors (for sufficiently large graphs) for any $\varepsilon > 0$. Choosing the graphs large enough and setting $\varepsilon' = \varepsilon/h$, we can bound the term $\binom{f(h)}{h}$ by $n^{\varepsilon' \cdot h} = n^\varepsilon$.

Theorem 8. *Let $\mathcal{G}$ be a nowhere-dense class and let H be a graph. For every $\varepsilon > 0$ there exists $N_\varepsilon \in \mathbf{N}$, such that for any graph $G \in \mathcal{G}, |G| > N_\varepsilon$ there exists an algorithm to count the appearances of H as a subgraph of G in time*

$$O\left(6^h \cdot h^{h+2} n^{1+\varepsilon}\right)$$

where f is a function depending only on the graph class.

Finally, we would like to point out that this counting algorithm is trivially parallelizable.

Measure	Definition	Localized
Closeness [90]	$c_C(v) = (\sum_{u \in V(G)} d(v,u))^{-1}$	$c^r_C(v) = (\sum_{u \in N^r(v)} d(v,u))^{-1}$
Harmonic [82]	$c_H(v) = \sum_{u \in V(G)} d(v,u)^{-1}$	$c^r_H(v) = \sum_{u \in N^r(v)} d(v,u)^{-1}$
Lin's index [60]	$c_L(v) = \frac{\lvert\{v \mid d(v,v) < \infty\}\rvert^2}{\sum_{u \in V(G): d(v,u) < \infty} d(v,u)}$	$c^r_L(v) = \frac{\lvert N^r[v]\rvert^2}{\sum_{u \in N^r[v]} d(v,u)}$

Table 3: Distance-based centrality measures with localized variants that can be computed in linear time on graphs of bounded expansion.

6.2 Localized Centrality

Centrality is a notion used to ascribe the relative importance of a vertex in the network. A centrality measure is a real-valued function that assigns each vertex of the network some value with the understanding that higher values correspond to more central vertices. Depending on the application, "central"vertices need not be those with high degree (for example, a cut-vertex may have high centrality as it is the only way for information to flow between two large subgraphs). There have been a wide variety of centrality scores introduced in the literature, including degree centrality, closeness centrality, eigenvector centrality, betweenness centrality and others. For a comprehensive introduction to centrality measures in social networks see, for instance, [33,34]. There are several recent articles devoted to the topic of centrality measures in general [52, 84]. In this section, we consider localized variants of measures similar to the *closeness centrality* introduced by Sabidussi [90] (see Table 3). The global versions of all these measures require one to compute the distance between all vertex pairs in the network, a sub-routine where the fastest known algorithm (due to Brandes [16]) is $O(n(n+m))$ which in the context of sparse networks reduces to quadratic time.

In these localized variants, we compute the measure of a vertex with respect to its r^{th} neighborhood rather than with respect to the whole graph. We give linear time algorithms for computing these measures on graphs of bounded expansion for every constant r. As the value of r increases, the value of the measure computed for a vertex approaches its unlocalized variant at the expense of an increase in running time.[3] The measures in question and their localized variants are listed in Table 3.

[3]For general graphs, we can compute these localized variants in time $O(n(n+m))$, by performing a breadth-first search from every vertex, for instance. We do not know whether a better running time is possible.

One natural question is the utility of localized variants (and their accuracy in reflecting the global measure). We remark that Marsden demonstrated that for some networks, calculating the measure for a vertex v inside its closed neighborhood $G[N[v]]$ can be used as a viable substitute for the full measure [62]. In the context of computer networks, Pantazopoulos et al. [83] consider local variants which lend themselves to distributed computing and found a close correlation to the full measures on a sample of networks. We can show experimentally that our variants reliably capture the top ten percent in (arbitrarily) selected networks of our real-world corpus. To that end, we compare the top 10 percent as identified by our localized variants to those 10 percent identified by the respective full centrality measure. Specifically, we use the Jaccard index [44]—defined as $|A \cap B|/|A \cup B|$ for two sets A, B—to measure similarity[4]

The results in Figure 3 suggest that already a value of r equal to about half the diameter yields very good results across all three measures. Note that we do not compare rankings; but rather only the difference between the sets of the identified top vertices—our experiments showed that rank ordering is not preserved reliably. Furthermore, there seems to be a slight positive tendency towards the localized version being better in larger networks, though it is hard to draw any conclusions on a small experiment like that.

While the localized Lin's index and the localized harmonic closeness work as depicted in Table 3, the localized closeness needs a small normalization tweak in order to yield good results: this is achieved by treating the $(r+1)$st neighborhood of very vertex as if it would contain all remaining vertices and adding this value accordingly (this obviously does not change the running time of the algorithm).

Let G be a graph from a class of bounded expansion and let $\vec{G}_r$ be the directed graph with in-degree bounded by $f(r)$, for some function f, that is obtained from G as by Proposition 4 (given at the beginning of Section 6. In the following, we let $N_r^-(v)$ denote the in-neighborhood of the vertex v in the directed graph $\vec{G}_r$. We assume that the vertices of G are ordered so that every vertex set has a unique representation as a tuple and, by slight abuse of notation, we use both representations interchangeably. For $v \in V(G)$ and $A = (a_1, a_2, \ldots, a_p) \subseteq N_r^-(v)$, define the *distance vector* from v to A as $\mathrm{dist}(v, A) := (\omega(a_1 v), \omega(a_2 v), \ldots, \omega(a_p v))$, where ω denotes the arc-labeling from Proposition 4. Since $a_i \in N_r^-(v)$, $\omega(a_i, v) = d_G(a_i, v)$.

Definition 16. *Let $v \in \vec{G}_r$, $\varnothing \neq X \subseteq N_r^-(v)$, $\alpha : V(G) \to \mathbf{R}$ a vertex*

[4]Since we compare sets of equal size, the measures precision and recall—and accordingly the F_1-score—are all the same. The Jaccard index is better suited for this situation.

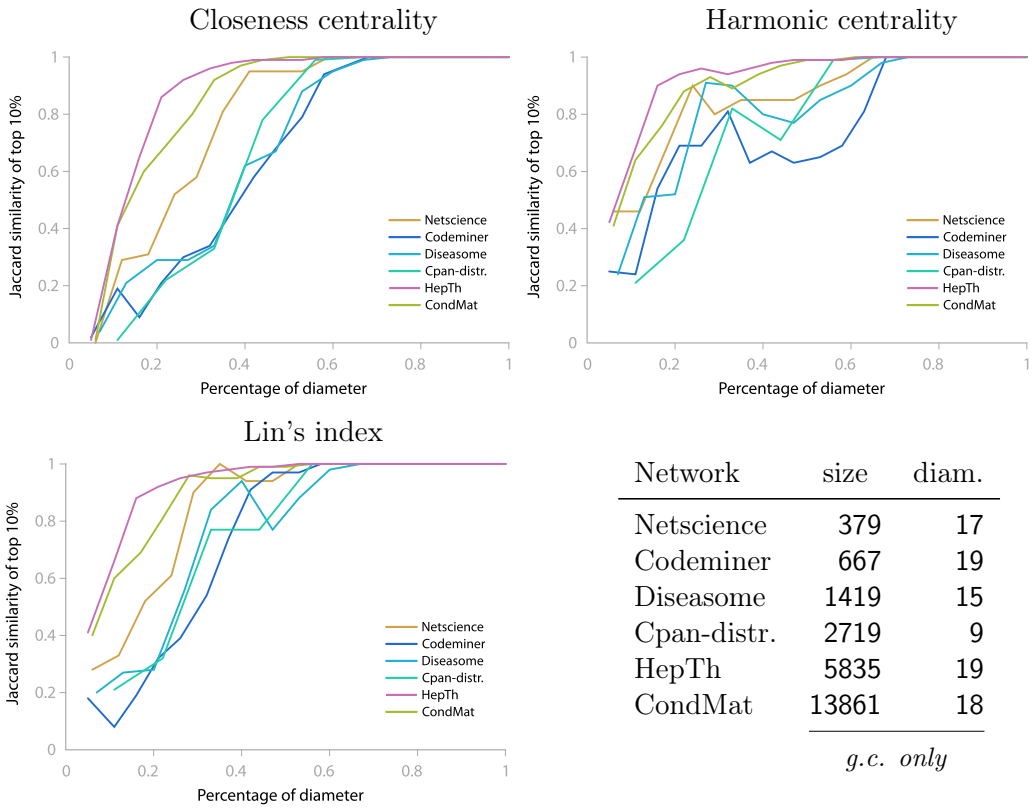

Network	size	diam.
Netscience	379	17
Codeminer	667	19
Diseasome	1419	15
Cpan-distr.	2719	9
HepTh	5835	19
CondMat	13861	18

g.c. only

Figure 3: Quality of localized centrality measures in terms of similarity of the top ten percent vertices against the full centrality measure. Measurements were taken only in the giant component of the networks displayed in the table.

weighting and let $\bar{d} \in [r]^{|X|}$ *be a distance vector. We define*

$$N(v, X, \bar{d}) := \{v \neq u \in V(G) \mid N_r^-(v) \cap N_r^-(u) = X \text{ and } dist(u, X) = \bar{d}\}$$

as those vertices whose in-neighborhood in $\vec{G}_r$ *overlap with the in-neighborhood of* v *in exactly* X *and whose distance-vector to* X *is exactly* $\bar{d}$*. Then the* query-function c_α *is defined as*

$$c_\alpha(v, X, \bar{d}) := \sum_{u \in N(v,X,\bar{d})} \alpha(u).$$

Lemma 18. *Given* $\vec{G}_r$*, one can compute a data structure in time* $O(n)$ *such that queries* $c_\alpha(v, X, \bar{d})$ *as in Definition 16 can be answered in constant time.*

Proof. We define an auxiliary dictionary R indexed by vertex sets $X \subseteq N^-r(v)$, for some vertex v. At each entry $v \in \vec{G}_r$, we will store another dictionary indexed by distance vectors which in turn stores a simple counter.

We initialize R as follows: for every $v \in \vec{G}_r, X \subseteq N_r^-(v)$ and every distance vector $\bar{d} \in [r]^{|X|}$, set $R[X][\bar{d}] = 0$. Note that in total, R contains only $O(n)$ entries since all in-neighborhoods in $\vec{G}_r$ have constant size. We can implement R as a hash-map to achieve the desired (expected) constant-time for insertion and look-up, though this would yield a randomized algorithm. A possible way to implement R on a RAM deterministically is the following: We store the key $X = \{x_1, x_2, \ldots, x_p\}$ at address $x_1 + n \cdot x_2 + \ldots + n^p \cdot x_p$. This uses addresses up to size n^c, for some constant c, but since we only insert $O(n)$ keys the setup takes only linear time. Our later queries to R will be restricted to keys that are guaranteed to be contained in the dictionary, thus we will never visit a register that has not been initialized.

Now, for every $v \in \vec{G}_r$, every $X \subseteq N_r^-(v)$, increment the counter $R[X][\text{dist}(v, X)]$ by $\alpha(v)$. We now claim that queries of the form $c_\alpha(v, X, \bar{d})$ can be computed using inclusion-exclusion as follows:

$$c_\alpha(v, X, \bar{d}) = \sum_{X \subseteq Y \subseteq N_r^-(v)} (-1)^{|Y \setminus X|} \sum_{\bar{d}' : \bar{d}'|_X = \bar{d}} R[Y][\bar{d}'].$$

The computation of the sum clearly takes constant time[5]. We now prove that it indeed computes the quantity $c_\alpha(v, X, \bar{d})$.

Consider a vertex $u \in \vec{G}_r$, such that $N_r^-(u) \cap N_r^-(v) = X$ and $\text{dist}(u, X) = \bar{d}$. We argue that $\alpha(u)$ is counted once by the above sum: $\alpha(u)$ is not counted by any $R[Y][\cdot]$ with $Y \supsetneq X$, therefore only the term where $X = Y$ counts $\alpha(u)$ and does so exactly once. It remains to be shown that the weight of vertices that do not conform with Definition 16 are either not counted by the sum or cancel out.

Consider a vertex $w \in \vec{G}_r$ with such that $\text{dist}(w, X) \neq \bar{d}$. The weight of such a vertex is not counted by the above sum, since $\alpha(w)$ is only counted in entries of R that do not occur as summands.

Finally, consider a vertex $w' \in \vec{G}_r$ with $N_r^-(w') \cap N^- r(v) = Z$ where $X \subsetneq Z \subseteq N_r^-(v)$ and such that $\text{dist}(w', X) = \bar{d}$. The weight of this vertex is counted in each term of

$$\sum_{X \subseteq Y \subseteq Z} (-1)^{|Y \setminus X|} R[Y][\text{dist}(w', Z)|_Y]$$

since

$$\sum_{X \subseteq Y \subseteq Z} (-1)^{|Y \setminus X|} = \sum_{0 \leqslant k \leqslant n} (-1)^k \binom{n}{k} = 0$$

we know that the signs cancel out and thus $\alpha(w')$ does not contribute to $c_\alpha(v, X, \bar{d})$. Hence the above sum computes exactly the query $c_\alpha(v, X, \bar{d})$. □

[5]We tacitly assume that the weights α only assign numbers polynomially bounded by the size of the graph.

Theorem 9. *Let $\mathcal{G}$ be a graph class of bounded expansion, $G \in \mathcal{G}$ a graph and $r \in \mathbf{N}$ an integer. Then one can compute the quantities $\alpha_d(v) = \sum_{w \in N^d(v)} \alpha(w)$ for all $v \in G, d \leqslant r$ in linear time.*

Proof. By Theorem 4 we can compute $\vec{G}_r$ in linear time, thus we can employ Lemma 18 to answer queries as defined in Definition 16 in constant time. To compute the quantity $\alpha_d(v)$ for all $0 < d \leqslant r$ and $v \in V(G)$, we proceed as follows. Initialize an array C by setting $C[v][d] = 0$ for every $v \in G, 1 < d \leqslant r$.

Now for every $v \in V(G)$, every $X \subseteq N_r^-(v)$ and every distance vector $\bar{d} \in [r]^{|X|}$, update C via

$$C[v][\min(\bar{d} + \operatorname{dist}(v, X))] \leftarrow C[v][\min(\bar{d} + \operatorname{dist}(v, X))] + c_\alpha(v, X, \bar{d})$$

and then apply the correction

$$C[v][\min(\operatorname{dist}(v, X) + \operatorname{dist}(v, X))] \leftarrow C[v][\min(\operatorname{dist}(v, X) + \operatorname{dist}(v, X))] - 1$$

in both cases with the convention that we dismiss entries where $\min(\bar{d} + \operatorname{dist}(v, X)) > r$. The second case corrects for the fact that the query $c_\alpha(v, N^-(v), \operatorname{dist}(v, N^-(v)))$ will count the vertex v itself.

At this point, $C[v][d]$ contains the sum of weights of vertices u for which $\min(\operatorname{dist}(v, X) + \operatorname{dist}(u, X)) = d$ where $X = N_r^-(v) \cap N_r^-(u) \neq \varnothing$. This follows directly from the definition of c_α.

By Theorem 4, every pair of vertices of distance $< r$ in G either is connected by an arc or they share a common in-neighbor in $\vec{G}_r$. Accordingly, we update the values of C as follows: for every $uv \in \vec{E}(\vec{G}_r)$

- if $N_r^-(u) \cap N_r^-(v) = \varnothing$, the weights of the vertices u and v were not counted in $C[v][\cdot], C[u][\cdot]$ respectively, thus we update C via

$$C[v][\omega(uv)] \leftarrow C[v][\omega(uv)] + \alpha(u)$$
$$C[u][\omega(uv)] \leftarrow C[u][\omega(uv)] + \alpha(v)$$

- if $X = N_r^-(u) \cap N_r^-(v) \neq \varnothing$, the weights of the vertices u and v were counted in $C[v][d']$ and $C[u][d']$ for $d' = \min(\operatorname{dist}(u, X) + \operatorname{dist}(v, X))$, respectively. Since d' might be larger than $\omega(uv)$ (but cannot be smaller), we update C via

$$C[v][d'] \leftarrow C[v][d'] - \alpha(u)$$
$$C[u][d'] \leftarrow C[u][d'] - \alpha(v)$$
$$C[v][\omega(uv)] \leftarrow C[v][\omega(uv)] + \alpha(u)$$
$$C[u][\omega(uv)] \leftarrow C[u][\omega(uv)] + \alpha(v)$$

  where we again ignore the update of $C[\cdot][d']$ if $d' > r$.

Note that this procedure is problematic if both uv and vu are present in the graph, since then the this correction would (wrongly) be applied twice. The simple solution is that in the case of both arcs being present we only apply the above update for that arc where the start vertex is smaller than the end vertex, i.e. to uv if $u < v$ and vu otherwise.
At this point, $C[v][d]$ contains the sum of weights of vertices u for which either

- the value $d = \min(\text{dist}(v, X) + \text{dist}(u, X))$ where $X = N_r^-(v) \cap N_r^-(u) \neq \varnothing$ and $uv \notin \vec{E}(\vec{G}_r)$,
- or $d = \omega(uv)$ and $uv \in \vec{E}(\vec{G}_r)$.

Thus by Theorem 4 we have that $C[v][d] = \alpha_d(v)$ for $d < r$ and $v \in G$. Since all of the above operations take time linear in $|V(G)|$, the claim follows. □

If we take $\alpha(\cdot) = 1$, the above algorithm counts exactly the sizes of the d^{th} neighborhoods of each vertex, for $d < r$. Thus it can be used to compute the r-centric centrality measures presented in Table 3.

Corollary 4. *Let $\mathcal{G}$ be a graph class of bounded expansion, $G \in \mathcal{G}$ a graph and $r \in \mathbf{N}$ an integer. Then the r-centric closeness, harmonic centrality and Lin's index can be computed for all vertices of G in total time $O(|V(G)|)$.*

7 Conclusion and Open Problems

We propose unifying structural graph algorithms with complex network analysis by searching for observable structural properties that satisfy the litmus test of enabling efficient algorithms for network analysis. We presented theoretical and empirical results that support our hypothesis that complex networks are structurally sparse in a well-defined and robust sense. Efficient algorithms are known for networks of bounded expansion [30, 40, 73], and we show that for key network analysis problems these algorithms can be even further improved. On the theoretical side, we show that several random graph models of complex networks exhibit bounded expansion with high probability, although not all do—suggesting an interesting dichotomy of networks. On the experimental side, we confirm these mathematical results, and show that many real-world complex networks additionally appear to exhibit bounded expansion as measured using specialized colorings. This new approach enables fast algorithms to analyze features including communities, centrality, and motifs while more broadly providing a rigorous framework for a deeper understanding of real-world networks and related models.

There are a plethora of random graph models specifically designed to mimic properties of complex networks. Which of these models exhibit structural sparsity (and which ones do not)? There is also room for debate about how to establish that a model will generate graphs with certain properties *in practice*. Asymptotic behavior is only a proxy, although we took care to provide details on the speed of convergence in our proofs where if possible. As exemplified by the relatively weak result about the Barabási-Albert model, the practical implications are sometimes difficult to judge.

On the algorithmic side, there are several key challenges remaining. Does there exist a better algorithm/heuristic to obtain low treedepth colorings, in particular taking into consideration the special structure of complex networks? Does a good coloring algorithm exist that provides a trade-off between the number of colors and the treedepth of subgraphs induced by few color classes? Can we compute or approximate lower bounds for either χ_p or $\widetilde{\nabla}_r$ with reasonable margins of error? Both would likely improve our current empirical understanding of the grad of networks. We would also like to investigate whether the grad for small depths is a reliable measure to differentiate networks; both our empirical and theoretical results seem to indicate so.

Finally, algorithms exploiting low grad should be tested extensively via computational experiments, to ascertain the feasibility of applying these techniques to real-world networks.

Acknowledgements.
E. Demaine supported in part by NSF grant CCF-1161626 and DARPA GRAPHS/AFOSR grant FA9550-12-1-0423.

Blair D. Sullivan supported in part by DARPA GRAPHS/SPAWAR Grant N66001-14-1-4063, the Gordon & Betty Moore Foundation, and the National Consortium for Data Science,

Peter Rossmanith, Fernando Sánchez Villaamil supported by DFG Project RO 927/13-1 "Pragmatic Parameterized Algorithms.".

Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of DARPA, SSC Pacific, AFOSR, the Moore Foundation, or the NCDS.

References

[1] CPAN explorer. http://cpan-explorer.org/. Accessed: July 2009.

[2] Datasets - gephi:wiki. https://wiki.gephi.org/index.php/Datasets. Accessed: 2015-01-27.

[3] L. A. Adamic and N. Glance. The political blogosphere and the 2004 US election. In *Proceedings of the WWW-2005 Workshop on the Weblogging Ecosystem*, 2005.

[4] A. Adcock, B. D. Sullivan, and M. Mahoney. Tree-like structure in large social and information networks. In *ICDM*, pages 1–10, 2013.

[5] I. Albert and R. Albert. Conserved network motifs allow protein-protein interaction prediction. *Bioinformatics*, 20(18):3346–3352, 2004.

[6] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, Jan. 2002.

[7] N. Alon. A note on network reliability. In *Discrete Probability and Algorithms*, pages 11–14. Springer, 1995.

[8] L. Babai, W. M. Kantor, and E. M. Luks. Computational complexity and the classification of finite simple groups. In *Proc. of 24th FOCS*, pages 162–171. IEEE, 1983.

[9] F. Ball, D. Sirl, and P. Trapman. Threshold behaviour and final outcome of an epidemic on a random network with household structure. *Adv. in Appl. Probab.*, 41(3):765–796, 09 2009.

[10] A. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[11] E. A. Bender and E. R. Canfield. The asymptotic number of labeled graphs with given degree sequences. *Jour. Comb. Theory, Series A*, 24(3):296–307, 1978.

[12] M. Biskup. On the scaling of the chemical distance in long-range percolation models. *Annals of Probability*, pages 2938–2977, 2004.

[13] B. Bollobás, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *Random Structures & Algorithms*, 31(1):3–122, 2007.

[14] B. Bollobás and O. Riordan. The diameter of a scale-free random graph. *Combinatorica*, 18(3):279–290, 2001.

[15] B. Bollobás, O. Riordan, J. Spencer, and G. Tusnády. The degree sequence of a scale-free random graph process. *Random Structures and Algorithms*, 24(1):5–34, 2004.

[16] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.

[17] T. Britton, M. Deijfen, and A. Martin-Löf. Generating simple random graphs with prescribed degree distribution. *Journal of Statistical Physics*, 124(6):1377–1397, 2006.

[18] D. Bu, Y. Zhao, L. Cai, H. Xue, X. Zhu, H. Lu, J. Zhang, S. Sun, L. Ling, N. Zhang, et al. Topological structure analysis of the protein–protein interaction network in budding yeast. *Nucleic acids research*, 31(9):2443–2450, 2003.

[19] J. Bu, C. Chen, X. Liu, M. Song, L. Zhang, and Q. Zhao. Probabilistic graphlet transfer for photo cropping. *Trans. Image Processing*, 22(2):802–815, 2013.

[20] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090. ACM, 2011.

[21] F. Chung and L. Lu. The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences*, 99(25):15879–15882, 2002.

[22] F. Chung and L. Lu. Connected components in random graphs with given expected degree sequences. *Annals of combinatorics*, 6(2):125–145, 2002.

[23] F. Chung and L. Lu. The small world phenomenon in hybrid power law graphs. In *Complex networks*, pages 89–104. Springer, 2004.

[24] A. Clauset, C. R. Shalizi, and M. E. Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[25] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on graphs of bounded-genus and H-minor-free graphs. *J. ACM*, 52(6):866–893, 2005.

[26] E. D. Demaine and M. Hajiaghayi. The bidimensionality theory and its algorithmic applications. *The Computer Journal*, 51(3):292–302, 2008.

[27] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Decomposition, approximation, and coloring of odd-minor-free graphs. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, pages 329–344, Austin, Texas, January 17–19 2010.

[28] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Contraction decomposition in H-minor-free graphs and algorithmic applications. In *Proc. of 43rd STOC*, pages 441–450, June 6–8 2011.

[29] Z. Dvořák. *Asymptotical Structure of Combinatorial Objects*. PhD thesis, Charles University, Faculty of Mathematics and Physics, 2007.

[30] Z. Dvořák, D. Kráľ, and R. Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36, 2013.

[31] M. Farrell, T. Goodrich, N. Lemons, F. Reidl, F. Sánchez Villaamil, and B. D. Sullivan. Hyperbolicity, degeneracy, and expansion of random intersection graphs. *arXiv preprint arXiv:1409.8196*, 2014.

[32] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Bidimensionality and kernels. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 503–510, Austin, Texas, 2010.

[33] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.

[34] L. C. Freeman. Centrality in social networks: Conceptual clarification. *Social Networks*, 1:215–239, 1978/79.

[35] S. Gago and D. Schlatter. Bounded expansion in web graphs. *Comment. Math. Univ. Carolin*, 50(2):181–190, 2009.

[36] J. Gajarský, P. Hlinený, J. Obdrzálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. Sánchez Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. In *Proc. of 21st ESA*, pages 529–540, 2013.

[37] Y. Gao. Treewidth of Erdős–Rényi random graphs, random intersection graphs, and scale-free random graphs. *Disc. Appl. Math.*, 160(4):566–578, 2012.

[38] M. Girvan and M. E. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.

[39] K.-I. Goh, M. E. Cusick, D. Valle, B. Childs, M. Vidal, and A.-L. Barabási. The human disease network. *Proceedings of the National Academy of Sciences*, 104(21):8685–8690, 2007.

[40] M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 89–98. ACM, 2014.

[41] W. Hayes, K. Sun, and N. Pržulj. Graphlet-based measures are suitable for biological network comparison. *Bioinformatics*, 29(4):483–491, 2013.

[42] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.

[43] S. Itzkovitz, R. Levitt, N. Kashtan, R. Milo, M. Itzkovitz, and U. Alon. Coarse-graining and self-dissimilarity of complex networks. *Physical Review E*, 71(1):016127, 2005.

[44] P. Jaccard. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bulletin del la Société Vaudoise des Sciences Naturelles*, 37:547–579, 1901.

[45] S. Janson. The probability that a random multigraph is simple. II. *Journal of Applied Probability*, 51:123–137, 2014.

[46] T. Jiang. Compact topological minors in graphs. *Journal of Graph Theory*, 67(2):139–152, 2011.

[47] E. J. Kim, A. Langer, C. Paul, F. Reidl, P. Rossmanith, I. Sau, and S. Sikdar. Linear kernels and single-exponential algorithms via protrusion-decompositions. In *Proc. of 40th ICALP*, pages 613–624, 2013.

[48] J. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845–845, 2000.

[49] J. Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proc. of 32nd STOC*, STOC '00, pages 163–170. ACM, 2000.

[50] B. Klimt and Y. Yang. Introducing the enron corpus. In *CEAS*, 2004.

[51] D. E. Knuth. Stanford GraphBase: A platform for combinatorial algorithms. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 41–43, 1993.

[52] D. Koschützki, K. A. Lehmann, L. Peeters, S. Richter, D. Tenfelde-Podehl, and O. Zlotowski. Centrality indices. In *Network analysis*, volume 3418 of *Lecture Notes in Computer Science*, pages 16–61. Springer, 2005.

[53] V. Krebs. V. Krebs' Website. http://www.orgnet.com/. Accessed: 2015-01-26.

[54] M. Krivelevich, D. Reichman, and W. Samotij. Smoothed Analysis on Connected Graphs. In *Proc. of APPROX/RANDOM 2014*, volume 28 of *LIPIcs*, pages 810–825, 2014.

[55] O. Kuchaiev, T. Milenković, V. Memišević, W. Hayes, and N. Pržulj. Topological network alignment uncovers biological function and phylogeny. *J. R. Soc. Interface*, 7(50):1341–1354, 2010.

[56] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2, 2007.

[57] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/data, June 2014.

[58] J. Leskovec, K. Lang, A. Dasgupta, and M. Mahoney. Statistical properties of community structure in large social and information networks. In *WWW '08: Proceedings of the 17th International Conference on World Wide Web*, pages 695–704, 2008.

[59] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009.

[60] N. Lin. *Foundations of Social Research*. McGraw-Hill New York, 1976.

[61] D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and S. M. Dawson. The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations. *Behavioral Ecology and Sociobiology*, 54(4):396–405, 2003.

[62] P. V. Marsden. Egocentric and sociocentric measures of network centrality. *Social Networks*, 24(4):407–422, 2002.

[63] A. Masoudi-Nejad, F. Schreiber, and Z. R. M. Kashani. Building blocks of biological networks: a review on major network motif discovery algorithms. *IET Systems Biology*, 6:164–174, 2012.

[64] R. Mathon. A note on the graph isomorphism counting problem. *Inform. Proc. Lett.*, 8(3):131–136, 1979.

[65] T. Milenković, V. Memišević, A. K. Ganesan, and N. Pržulj. Systems-level cancer gene identification from protein interaction network topology applied to melanogenesis-related functional genomics data. *J. R. Soc. Interface*, 7(44):423–437, 2010.

[66] T. Milenković and N. Pržulj. Uncovering biological network function via graphlet degree signatures. *Cancer Informatics*, 6(6):257, 2008.

[67] S. Milgram. The small-world problem. *Psychology Today*, 61(1), 1967.

[68] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

[69] M. Molloy and B. A. Reed. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms*, 6(2/3):161–180, 1995.

[70] M. Molloy and B. A. Reed. The size of the giant component of a random graph with a given degree sequence. *Combin., Probab. Comput.*, 7(3):295–305, 1998.

[71] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion I. decompositions. *Eur. Jour. Comb.*, 29(3):760–776, 2008.

[72] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion II. algorithmic aspects. *Eur. Jour. Comb.*, 29(3):777–791, 2008.

[73] J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms & Combinatorics*. Springer, 2012.

[74] J. Nešetřil, P. Ossona de Mendez, and D. R. Wood. Characterisations and examples of graph classes with bounded expansion. *Eur. Jour. Comb.*, 33(3):350–373, 2012.

[75] M. Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.

[76] M. E. Newman. Newman's network data. http://www-personal.umich.edu/~mejn/netdata/. Accessed: 2015-01-26.

[77] M. E. Newman. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2):404–409, 2001.

[78] M. E. Newman. Random graphs as models of networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook of Graphs and Networks*. Wiley-VCH, 2003.

[79] M. E. Newman. The structure and function of complex networks. *SIAM Reviews*, 45(2):167–256, 2003.

[80] M. E. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical review E*, 74(3):036104, 2006.

[81] M. E. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64(2), 2001.

[82] T. Opsahl, F. Agneessens, and J. Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3):245–251, 2010.

[83] P. Pantazopoulos, M. Karaliopoulos, and I. Stavrakakis. On the local approximations of node centrality in internet router-level topologies. In *Proc. of 7th IFIP IWSOS*, pages 1–12, 2013.

[84] N. Perra and S. Fortunato. Spectral centrality measures in complex networks. *Physical Review E*, 78, 2008.

[85] P. Pongpaibool, S. Pukkawanna, and V. Visoottiviseth. Lightweight detection of DoS attacks. In *Proc. of 15th ICON*, pages 77–82. IEEE, 2007.

[86] N. Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):177–183, 2007.

[87] F. Reidl, P. Rossmanith, F. Sánchez Villaamil, and S. Sikdar. A faster parameterized algorithm for treedepth. In *Automata, Languages, and Programming*, pages 931–942. Springer, 2014.

[88] P. Ribeiro, F. Silva, and M. Kaiser. Strategies for network motifs discovery. In *Proc. of 5th E-Science*, pages 80–87. IEEE, 2009.

[89] M. Ripeanu, A. Iamnitchi, and I. Foster. Mapping the Gnutella network. *IEEE internet computing*, 6(1):50–57, 2002.

[90] G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966.

[91] O. Sporns and R. Kötter. Motifs in brain networks. *PLoS Biology*, 2(11):369, 2004.

[92] J. Stehlé, N. Voirin, A. Barrat, C. Cattuto, L. Isella, J. Pinton, M. Quaggiotto, W. Van den Broeck, C. Régis, B. Lina, and P. Vanhems. High-resolution measurements of face-to-face contact patterns in a primary school. *PLOS ONE*, 6(8):e23176, 08 2011.

[93] S. Takemura, A. Bharioke, Z. Lu, A. Nern, S. Vitaladevuni, P. K. Rivlin, W. T. Katz, D. J. Olbris, S. M. Plaza, P. Winston, T. Zhao, J. A. Horne, R. D. Fetter, S. Takemura, K. Blazek, L.-A. Chang, O. Ogundeyi, M. A. Saunders, V. Shapiro, C. Sigmund, G. M. Rubin, L. K. Scheffer, I. A. Meinertzhagen, and D. B. Chklovskii. A visual motion detection circuit suggested by Drosophila connectomics. *Nature*, 500(7461):175–181, Aug. 2013.

[94] S. Teso, J. Staiano, B. Lepri, A. Passerini, and F. Pianesi. Ego-centric graphlets for personality and affective states recognition. In *Proc. of SocialCom 2013*, pages 874–877. IEEE, 2013.

[95] J. Ugander, L. Backstrom, and J. Kleinberg. Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections. In *Proc. of the 22nd WWW*, pages 1307–1318. IW3C2, 2013.

[96] Y. J. Wang and G. Y. Wong. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*, 82(397):8–19, 1987.

[97] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

[98] W. W. Zachary. An information flow model for conflict and fission in small groups. *Journal of anthropological research*, pages 452–473, 1977.

[99] T. Zhou, G. Yan, and B.-H. Wang. Maximal planar networks with large clustering coefficient and power-law degree distribution. *Physical Review E.*, 71(4):046141, 2005.

Appendix

Clustering and bounded expansion

One could wonder if having a graph class with constant clustering coefficient is an impediment to the graph class having bounded expansion. This does not have to be the case, as exemplified by our result on the configuration model with households and the *random appolonian networks* model, which generates planar networks with a power-law degree distribution and constant clustering coefficient in the limit [99].

It is easy to show something even stronger for graph classes with bounded expansion, namely that we can "force" any constant coefficient. Note that if G is a graph with bounded average degree d, we can obtain a constant clustering as follows: for every vertex in the graph, add d^2 edges to its neighborhood. To see that this results in a constant clustering, note that every vertex with degree at most d now has a clique as its neighborhood—thus in the resulting graph G', the local clustering of such a vertex is necessarily equal to 1. Since the number of vertices that have degree at most the total average degree makes up a constant fraction of the whole vertex set, the resulting clustering coefficient will be bounded from below by a constant. However, we need to show that such an operation will not destroy the property of having bounded expansion.

Lemma 19. *Let $\mathcal{G}$ be a graph class. For any constant t, let $\mathcal{C}_t(\mathcal{G})$ be the class of graphs obtained from any graph of $\mathcal{G}$ by adding up to t edges to the neighborhood of every vertex. Then the grad of $\mathcal{C}_t(\mathcal{G})$ depends only on the grad of $\mathcal{G}$ and t. In particular, if $\mathcal{G}$ has bounded expansion, so does $\mathcal{C}_t(\mathcal{G})$.*

Proof. Take a graph $G \in \mathcal{G}$ and for each $v \in G$, let E_v be a set of at most t edges to be added to the neighborhood of v. Let G' be the graph obtained from G by adding the edge set $\bigcup_{v \in G} E_v$. We compute G' only increasing the grad by some function depending on the old grad and some constants.

Let $d = 2\nabla_0(G)$ denote the degeneracy of G. We compute a proper coloring of G with $d+1$ colors and obtain a partition $V_1, \ldots, V_{d+1}$ of the vertex set. This also naturally partitions the edge set $\bigcup_{v \in G} E_v$ into sets $E_1, \ldots, E_{d+1}$, where $E_i = \bigcup_{v \in V_i} E_v$ for $1 \leqslant i \leqslant d+1$.

Let us demonstrate how the edges from a single set E_i can be added to G through operations that do not increase the grad too much; this sequence of operations can then be simply executed for every $1 \leqslant i \leqslant d+1$ to obtain enough edges for the graph to have constant clustering. Fix a set E_i in the following. We first take the lexicographic product $G \bullet K_{t+1}$. For each vertex $v \in G$, the set $C_v := \{(v, u) \mid u \in K_{t+1}\}$ denotes the 'copies' of v in $G \bullet K_{t+1}$.

By Lemma 5.1 and Lemma 5.2 in [71], this increases the grad of the graph only by some function of t and the old grad. Now, consider some vertex $v \in V_i$ in whose neighborhood the edge set $E_v \subseteq E_i$ should be added. To this end, for every edge $ab \in E_v$ we take some vertex $v' \in C_v$ and contract it into an edge between C_a and C_b (by the choice of v', those vertices are exactly determined.) Since $|E_v| \leqslant t$, we can do this without exhausting C_v. Furthermore, since the vertices of V_i are independent, no vertex will both create *and* receive an edge. Therefore, after contracting each sets $C_v, v \in G$ in the remaining graph—i.e. excluding the vertices that wzere contracted into edges—we obtain exactly the graph G with the edges E_i added to it. This is, by construction, a 1-shallow minor of $G \bullet K_{t+1}$. Thus this last operation cannot increase the grad arbitrarily, but only by a function of t.

Applying the above steps of operations for each edge set E_i yields the graph G with the edges $\bigcup_{v \in G} E_v$ added to it. Since d is bounded by the grad of G, the resulting graph's grad only depends on t and the grad of G. □

We therefore conclude that clustering is not at all orthogonal to bounded expansion.